Supplement to:

How to Save a Leaky Ship: Capability Traps and the Failure of Win-Win Investments in Sustainability and Social Responsibility

John Lyneis
PA Consulting
10 Canal Park
Cambridge, MA 02141
(617) 225-2700
john.lyneis@paconsulting.com

John Sterman*
MIT Sloan School of Management
100 Main Street
Cambridge, MA 02139
(617) 253-1951
jsterman@mit.edu

* Corresponding author

Table of Contents

S1. Data Sources and Methods

As outlined in the main text, we employ a mixed methods approach that includes ethnographic interviews, statistical analysis of building maintenance and energy data, and simulation modeling. This supplement provides additional detail on data sources used, qualitative and quantitative research methods, and equation formulations for the simulation model. The final section describes how to simulate the model to replicate results presented in the paper.

The major data sources used are the following:

1. **Ethnographic interviews**: Thirty interviews conducted with maintenance mechanics, managers, engineers, representatives for academic departments, and MIT administrators. (Described in more detail in section S3 below);
2. **Database of campus building systems**: Database listing all major building systems, value, lifetime, and year of required capital renewal. Used to model the condition of building systems the effect of investment in capital renewal. (Described in more detail in section S2 below);
3. **Data on building energy usage**: Yearly energy usage for 111 campus buildings, three energy carriers (chilled water, steam and electricity), over a seven year period (2000-2006). Used to estimate the relationship between building aging and energy usage (described in section S3 below);
4. **Data on maintenance work orders**: Weekly data on maintenance work orders between 2005 and 2008. Data include work orders opened and closed, hours charged per work order, trade/type of work, and whether the work was reactive or preventive maintenance. The data are used to estimate parameters in the simulation model in two areas: (1) the allocation of maintenance time between reactive and preventive maintenance; and (2) the productivity of the maintenance workforce (work orders/person-hour), including the relationship between productivity and time pressure to complete work. Details on model formulations are included in section S5.

S2. Database of building systems

The model is based on a detailed engineering assessment of campus condition completed in 2007. MIT hired an external engineering services firm to inspect all campus buildings, document current condition and defects, and estimate the lifetime and renewal costs for all major building systems. The resulting quantitative database of findings is used to parameterize the simulation model. The database includes approximately 6600 items, representing all major building systems in every campus facility. Table S1 illustrates the level of detail. Key attributes of each item include:

- **Name**: The name includes details on the type of building system. The model includes each individual item in the engineering database. We group each of the roughly 6600 items into six categories corresponding to a standard industry classification scheme: exterior structures, interior structures, plumbing, electrical, HVAC, and other. The same grouping is applied to maintenance work order data to link the condition of buildings and building systems to the generation of defects that lead to maintenance work orders. Following the database, the model represents each item as being in "good condition" until the date at which the database indicates it will reach the end of its useful life and require renovation or replacement, at which point that item is moved from the "good condition" state to "needs renewal." When renewal is completed, the item moves back into the "good condition" state.

- **Cost**: The cost of renewing an item. In the simulation, the cost is the amount of capital investment required for an item to move from the "needs renewal" category to the "good condition" category.
- **Year**: The year is the year when renewal will be required. At the time of the assessment in 2007, numerous items were already in a "needs renewal" state. The sum of the value of these items is the deferred maintenance backlog in the model.
- **Lifetime**: The number of years an item can be in service before it requires renewal.
- **Building Number**: Campus building where system resides.
- **Asset Size**: Size in gross square feet (gsf) of the building where the system resides.

Item	ID	Name	Name Coded	Cost	Year	Type	Lifetime	Building Number	Asset Size	Category
3938	2588456	D20-Plumbing	Plumbing	1920	2013	REN	12			D2
3939	2622221	D20-Plumbing	Plumbing	6392.29	2005	REN	12			D2
3940	2649023	D20-Plumbing	Plumbing	192310	2010	REN	40			D2
3941	2649048	D20-Plumbing	Plumbing	192773	2023	REN	40			D2
3942	2527126	D3011-Oil Supply System	HVAC	112185	2029	REN	35			D3
3943	2534518	D3011-Oil Supply System	HVAC	59101	2029	REN	35			D3
3944	2466008	D3012-Gas Supply System	HVAC	4982	2020	REN	40			D3
3945	2466062	D3012-Gas Supply System	HVAC	4758	2019	REN	40			D3
3946	2466090	D3012-Gas Supply System	HVAC	1991076	2015	REN	40			D3
3947	2466116	D3012-Gas Supply System	HVAC	755254	2011	REN	40			D3
3948	2466143	D3012-Gas Supply System	HVAC	3782	2014	REN	40			D3
3949	2466172	D3012-Gas Supply System	HVAC	5461	2022	REN	40			D3
3950	2545045	D3012-Gas Supply System	HVAC	35732	2014	REN	40			D3
3951	2551751	D3012-Gas Supply System	HVAC	2058.36	2005	REN	40			D3
3952	2563272	D3012-Gas Supply System	HVAC	18136	2019	REN	40			D3
3953	2574853	D3012-Gas Supply System	HVAC	16544	2017	REN	40			D3
3954	2574979	D3012-Gas Supply System	HVAC	15092	2015	REN	40			D3
3955	2575124	D3012-Gas Supply System	HVAC	4340	2017	REN	40			D3
3956	2622296	D3012-Gas Supply System	HVAC	22913.7	2005	REN	40			D3
3957	2459725	D3020-Heat Generating Syst	HVAC	128904	2024	REN	30			D3
3958	2464788	D3020-Heat Generating Syst	HVAC	299703	2029	REN	30			D3
3959	2554644	D3020-Heat Generating Syst	HVAC	626095	2027	REN	30	Redacted		D3
3960	2556929	D3020-Heat Generating Syst	HVAC	181979	2022	REN	30			D3
3961	2562336	D3020-Heat Generating Syst	HVAC	142270	2025	REN	30			D3
3962	2563763	D3020-Heat Generating Syst	HVAC	1411009	2016	REN	30			D3
3963	2564571	D3020-Heat Generating Syst	HVAC	118361	2017	REN	30			D3
3964	2569527	D3020-Heat Generating Syst	HVAC	224897	2027	REN	30			D3
3965	2570702	D3020-Heat Generating Syst	HVAC	303745	2025	REN	30			D3
3966	2575094	D3020-Heat Generating Syst	HVAC	1367915	2013	REN	30			D3
3967	2581436	D3020-Heat Generating Syst	HVAC	354866	2027	REN	30			D3
3968	2582283	D3020-Heat Generating Syst	HVAC	126392	2027	REN	30			D3
3969	2582440	D3020-Heat Generating Syst	HVAC	55293.97	2005	REN	30			D3
3970	2582630	D3020-Heat Generating Syst	HVAC	120718	2026	REN	30			D3
3971	2582738	D3020-Heat Generating Syst	HVAC	230598	2025	REN	30			D3
3972	2588193	D3020-Heat Generating Syst	HVAC	96167	2025	REN	30			D3
3973	2622207	D3020-Heat Generating Syst	HVAC	1144795	2007	REN	30			D3
3974	46979	D3021-Boilers	HVAC	137457	2022	REN	20			D3
3975	441383	D3021-Boilers	HVAC	9211	2010	REN	40			D3
3976	2307183	D3030-Cooling Generating S	HVAC	7008	2017	REN	15			D3

Table S1: Illustrative example from the database of building systems

The database includes only items with a renewal year on or before 2030. As a result it is right censored. That is, it omits systems with long life spans in newer buildings, and systems in older buildings that have recently been renewed. For example, a system with a life span of 35 years or longer in a building built or renovated after 1995 would not appear in the database.

Figure S1 confirms the absence of longer lifetime items in newer buildings. The figure shows the average renewal cost of items per gsf against item lifetime for two groups of buildings: buildings built before 1980, and buildings built after 1980. In older buildings (some are as old as 100 years) the distribution includes items with lifecycles of 80 years or more. (These are typically exterior structures and not HVAC or electrical systems). In contrast, newer buildings are missing these longer lifecycle items.

Because we run simulations only through 2030, the impact of omitted items on simulation results is likely minimal. All building systems that would come up for renewal and contribute to increased maintenance defects during the simulated time frame are represented within the database. Nevertheless, because "good condition" items also produce defects that generate maintenance work orders and affect energy use (albeit at a lower rate), we must correct for the right-censored omission of these systems from the database so that the proportion of items in good condition compared to those requiring renewal is accurately reflected.

To correct for the right-censoring in the database, we add items to the good condition stock based on the distribution shown in S1. We assume that all buildings, if fully represented, would show a distribution equivalent to the average of all pre-1980 buildings. Although some items in older buildings may have been renewed already and thus might also be omitted, we assume that the database is complete for old buildings, a fair assumption given the relative low rate of investment in renewal over the past few decades. We calculate the square footage of all newer buildings where items may be omitted. For each system lifetime, additional "good condition" items are introduced into the newer buildings assuming the same number and cost per gsf as in older buildings.

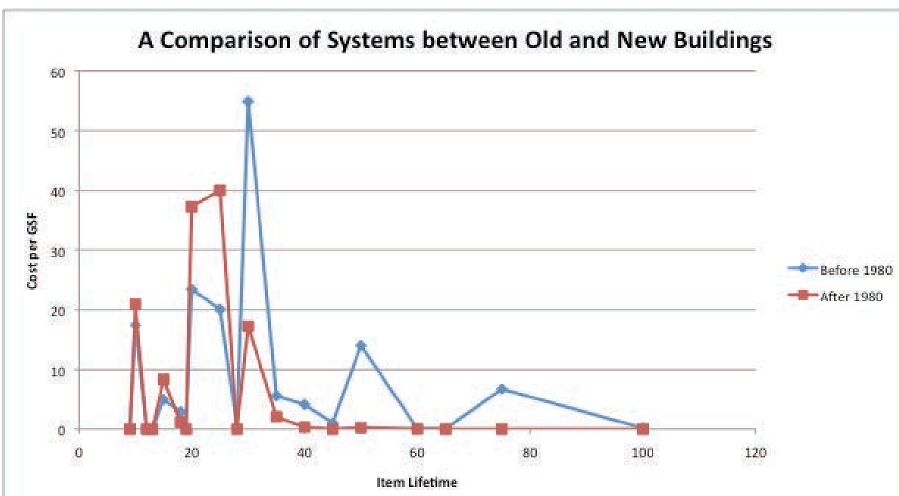

Figure S1: Comparison of Systems between Old and New Buildings

Each system contained in the database is represented individually in the model, with its associated cost, renewal date, and lifecycle. When an individual item reaches the end of its lifetime as specified in the database it changes state from "good condition" to "needs renewal." Items needing renewal remain in that state until sufficient capital investment is made for that item to return it to good condition. Items that are renewed are assigned a new future renewal date based on the lifecycle and time of renewal. More details on model formulation are provided in section S5 below.

S3. Qualitative Research Methods

Ethnographic interviews were important in both theory development and the formulation of the simulation model. The study began in 2007. At that time, advocates within the department of facilities and the broader MIT community believed that maintenance represented an opportunity for quick wins in reducing MIT's energy consumption and carbon footprint. To test this idea, we originally scheduled 2-3 interviews with managers in the maintenance department. These interviews identified a set of broader issues related to deferred maintenance and campus renewal. It was apparent early on that the maintenance organization was under some strain, and that these organizational dynamics were important to understanding why low-hanging fruit in the area of efficiency and proactive maintenance remained unpicked.

To investigate further, we conducted a total of 30 semi-structured interviews with a range of individuals in the maintenance organization, the MIT administration, and the broader community. Interviews were based around a questions designed to understand the daily pressures that individuals faced related to building maintenance, how these pressures had changed over time, what their origins were, and what if anything had been done to alleviate them. Sample interview questions are listed in Table S2:

Sample Questions used in Semi-Structured Interviews
Hourly Workers, Supervisors: - What are the top priorities communicated to you by management? - Are you encouraged to complete jobs quickly? How quickly? Do you ever feel under time pressure, or are you encouraged to take your time to do things right? - How often do you have to drop what you're doing to fix an emergency? What are typical emergencies? - What are typical preventive maintenance tasks? Do you often find problems that required fixing? Are these followed through on? - How often do you place your own work orders? - In your opinion, what is required to keep the systems that you're working on in good order? Is this done? - Do you have a good feel for the systems you work on? - How important are performance metrics to your job? Do they affect the work? - What are the biggest differences now compared to when you first started working at MIT? Five years ago? - Do you feel valued by management, by the MIT community as a whole? Managers, Administration: - What goals do you have for maintenance and campus renewal? Have those goals been met? What are the largest obstacles to meeting those goals? - In general, how tight are resources within facilities? How are budgeting decisions made? What happens when something happens that is not budgeted for? - How have R&M, and MIT facilities generally, changed since you arrived at MIT? What were the sources of those changes?

Table S2: Sample questions used in semi-structured Interviews

The specific questions asked in the interviews varied slightly depending upon the position of the individual interviewed. Table S3 lists several of the types of individuals interviewed together with the types of questions raised with each:

Interviewee(s)	**Description of Position**	**Interview Topics**
Maintenance Mechanics	Hourly maintenance workers (e.g. plumbers, electricians, HVAC mechanics)	Daily pressures, types of work performed, condition of equipment, priorities of management, institutional history
Maintenance Supervisors	Immediate supervisors of mechanics, generally assigned to a specific trade and/or campus zone	Daily pressures, priorities of management, workforce issues
Maintenance Managers	Managers of the maintenance department, responsible for overall performance	Departmental goals, policies, and challenges, metrics employed, priorities of senior administration
Facilities Engineers	Engineers responsible for supporting maintenance work	Condition of equipment, institutional history
Facilities Sustainability Managers	Individuals within the department of facilities responsible for energy efficiency initiatives	Sustainability initiatives, profile of energy consumption, contribution of maintenance and renewal to energy consumption
Departmental Facilities Liaisons	Individuals employed by academic departments who work with facilities to ensure that buildings meet department needs	Perspectives on performance of maintenance organization, departmental priorities, daily pressures, institutional history
Finance Administrators	Senior MIT administrators responsible for overall priorities and budgeting	Budgeting processes, goals and challenges related to campus renewal
Managers for new construction and capital renewal	Facilities managers responsible for overseeing new construction and renovation projects	Processes for new construction and capital renewal, profile of new buildings on campus
Facilities Administrators	Senior managers responsible for overall department operations (including maintenance, new construction, utilities, custodial services, and other Institute operations)	Budgeting challenges, goals and achievements

Table S3: List of types of individuals interviewed and topics covered

Interviews were recorded and transcribed. Themes emerging from interviews were then grouped and used to inform the development of the causal model presented in the main text.

After an initial round of ethnographic interviews, follow up interviews were conducted as the quantitative model was developed and tested. In addition to the topics above, these interviews included requests for quantitative data, questions about data sources, and questions related to model formulations and parameters.

S4. Estimating the relationship between building aging and energy use

As described in the main text, we model the change in building energy requirements over time using an exponential specification:

$$\frac{dE_k}{dt} = \frac{(E_k^* - E_k)}{\tau_k} \tag{S1}$$

where E_k is energy consumption per gsf, indexed by energy carrier $k \in$ {steam, chilled water, electricity}, E_k^* is the asymptotic energy requirement per square foot, i.e., energy consumption per gsf in the long run, assuming no additional investments in building renewal or efficiency, and $\tau_k$is time constant governing how long it takes energy efficiency to deteriorate. Theory and experience suggest that, in the absence of efficiency upgrades, the energy requirements of buildings and systems will rise over time as equipment and structures age and deteriorate, and that this increase will be asymptotic. The exponential form of eq. (S1) is a common and simple specification with the desired properties. The solution to eq. (S1) is

$$E_k = E_k^* - (E_k^* - E_k^0)\exp(-t/\tau_k) \tag{S2}$$

where E_k^0 is initial energy use. Taking the logarithm of the time-derivative yields a linear-in-parameters form we can use to estimate E^* and τ_k,

$$\ln(\dot{E}_k) = \ln\left(\frac{E_k^* - E_k^0}{\tau_k}\right) - t/\tau_k \tag{S3}$$

where $\dot{E}_k$is the rate of change in energy consumption per gsf.

To estimate E* and τ_k we proceed in several steps. First, we obtain individual building estimates for $\dot{E}_k$ by running linear regressions for each building-energy carrier combination with time as an independent variable. Second, we use these estimates and equation (S3) above to estimate τ_k. To do so, we regress estimates for $\dot{E}_k$ against building age (*t*) for each energy carrier using log-linear regression. Estimated slopes can then be used to calculate τ_k. Third, we use the full panel regression across all campus buildings and seven years of data to estimate E_k^*. Given estimates for τ_k obtained in steps one and two, estimates for E_k^* can be calculated using results from the full panel regression. We next provide more details on each of these steps.

The first step is to obtain estimates for $\dot{E}_k$ for each building and energy carrier. We use the equation:

$$E_{k,b}^t = E_{k,b}^0 + \dot{E}_{k,b}t + aH_t + bC_t \tag{S4}$$

where *b* is the building, *k* is the energy carrier and H and C are heating and cooling degree days for the year. Time (*t*) is an independent variable with slope equal to $\dot{E}_{k,b}$. The panel data sample includes energy usage in yearly increments over a seven year period (2000-2006), for 101 buildings and three energy carriers.

Results yield statistically significant estimates for $\dot{E}_{k,b}$ for 15 of 60 buildings for chilled water, 31 of 101 buildings for electricity, and 21 of 77 buildings for steam. (The number of buildings differs between carriers because not all buildings consume chilled water and steam).

If the hypothesized exponential model in (S1) is a good fit, we expect to see lower estimates for $\dot{E}_{k,b}$ for older buildings. Figure S2 shows plots of estimates for $\dot{E}_{k,b}$ against building age since major renovation. Plots are shown both for significant estimates only, and for all estimates.

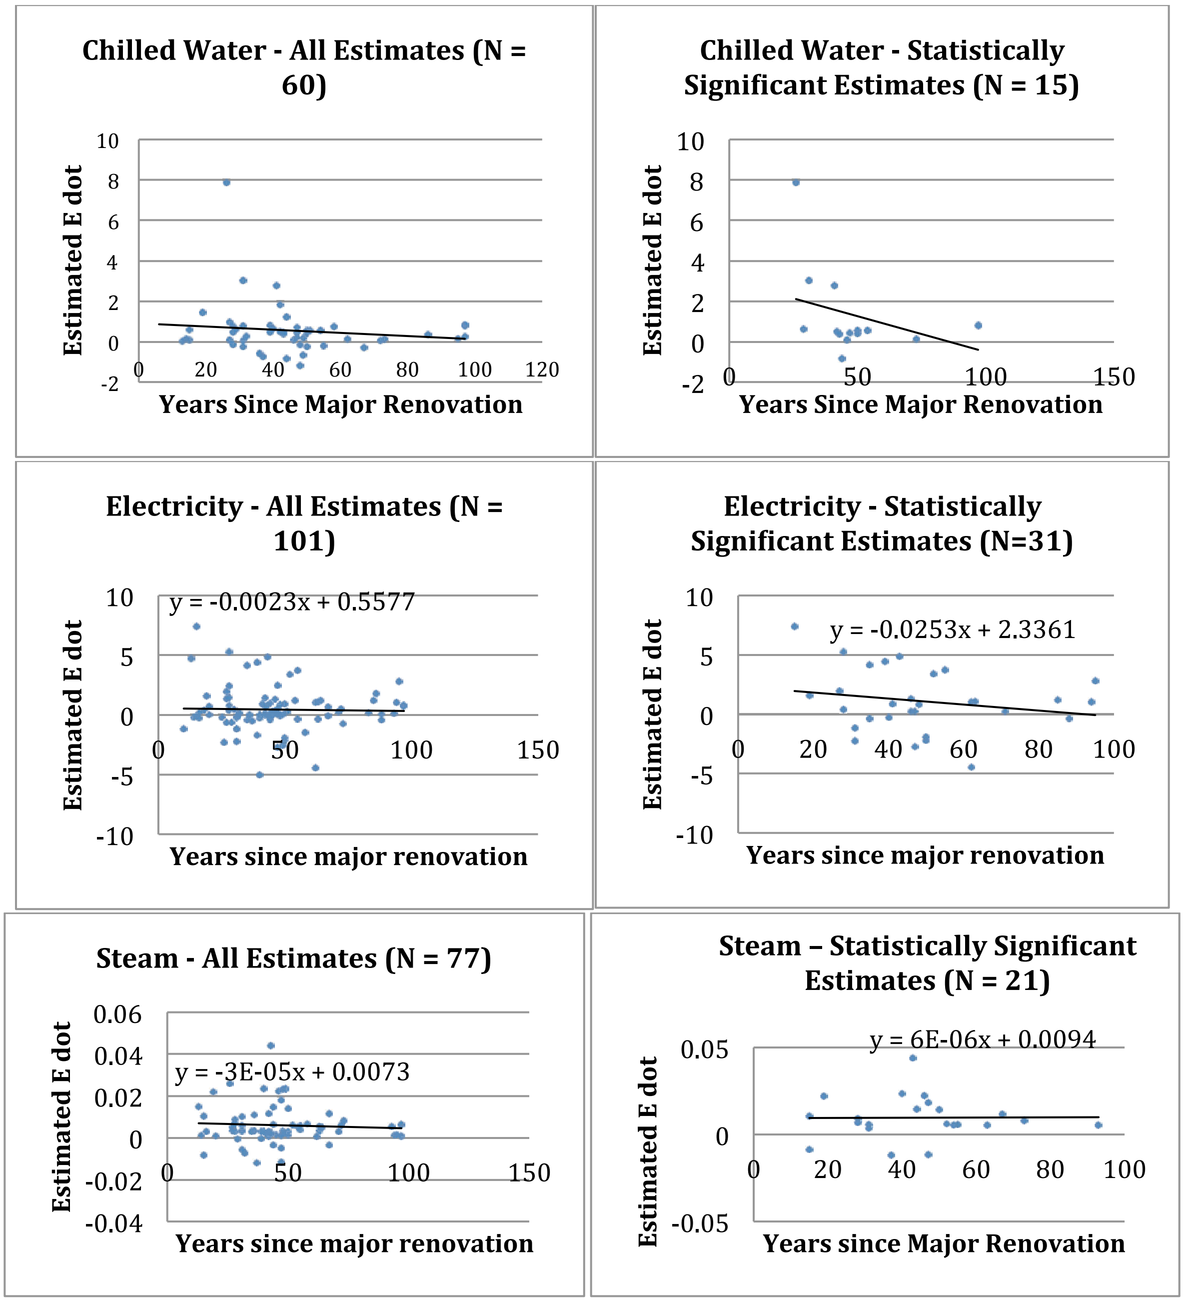

Figure S2: Plots of building age against regression estimates for $\dot{E}_k$

The second step uses information on building age together with individual building $\dot{E}_{k,b}$ estimates to estimate a single τ_k for each carrier. Equation (S3) gives the functional relationship between age (*t*) and $\dot{E}_{k,b}$. We regress $\dot{E}_{k,b}$ against building age, and use the estimate for the slope to calculate τ_k. Specifically, τ_k is equal to -1 divided by the estimated slope. The results are directionally as theory predicts but statistically significant only for steam (Table S4).

	Slope Estimate (standard error)	$\tau_k =$ -1/slope
Chilled Water	-0.00451 (.00672)	222.2 years
Electricity	-0.000584 (.0087)	1712 years
Steam	-0.00993* (0.00576)	100.7 years

Table S4: Regression Results for the Model in eq. S3.

Given the failure to obtain statistically significant results for chilled water and electricity, we attempt a second approach. We place buildings into 5 "buckets" based on their age, using 20-year increments, and run a separate panel regression to determine an estimate of $\dot{E}_k$ for each age bucket. We then repeat the analysis above. $\dot{E}_k$ estimates are shown in Table S5, and plotted in Figure S3. Regressing log($\dot{E}_k$) against age again gives estimates for τ of 37 years for chilled water, 552 years for electricity, and 117 years for steam. Due to the small number of data points in each bucket, however, these estimates are not statistically significant.

	Chilled Water		Electricity		Steam	
Age Bucket	E_k^* Estimate	N	E_k^* Estimate	N	E_k^* Estimate	N
0-20	.579* (.16)	9	1.54* (.4)	16	.007* (.003)	11
21-40	.91* (.24)	17	.21 (.24)	30	.0046* (.0013)	22
41-60	.39* (.15)	23	.4* (.15)	34	.008* (.0013)	29
61-80	.005 (.11)	4	-.07 (.47)	10	.0045* (.0009)	9
81-100	.53* (.24)	6	.8 *(.15)	12	.003* (.0009)	7

Table S5: Estimates of E_k^* for age buckets (smaller Ns for CW and Steam are due to the fact that not all buildings use CW and steam).

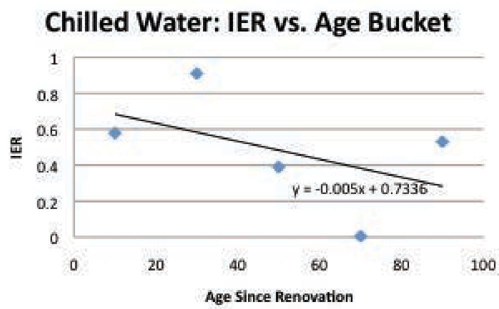

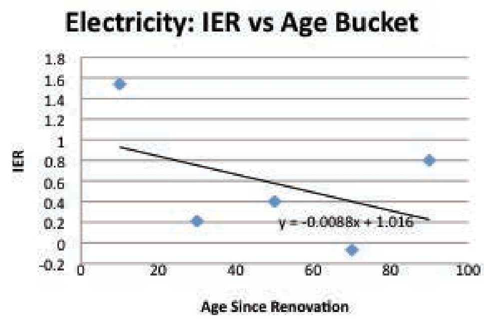

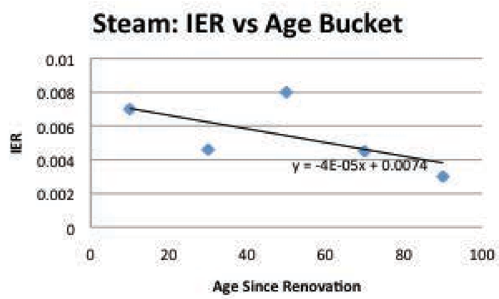

Figure S3: Plots of Building Age against the Estimated Increase in Energy Requirements (IER) for age buckets.

Considering the estimates from the two approaches, we note the following:

(i) The estimated deterioration time for chilled water in Table S3 is implausibly long. The chilled water system includes a relatively long-lived distribution system, but most of the value consists of the chiller plants themselves, which involve complex equipment including the motors, compressors, heat exchangers, controls and other elements of the refrigeration units. Rotating equipment, pressurized components and controls are not expected to last as long as the steam system, which is mechanically simpler. We therefore use the estimate from the bucket approach.

(ii) The estimated time constant for electricity from both methods indicates that there is no discernible slowdown in the rate at which electricity consumption per gsf rose over the time frame spanned by the data. The very long estimate arises from the fact that the data confound two processes, both of which are picked up by the time trend term in the regressions: Electricity use includes both the impact of aging and rising plug loads as the density of electronics and other equipment in offices, labs and dorms has risen, a process co-linear with building aging. We use the estimate from the bucket approach in the model, noting that the change electricity demand over time is essentially linear in the estimation period and the difference between the two methods of estimating the parameters of eq. S2 yield essentially the same results over the model time horizon (through 2030).

(iii) The estimates for steam are similar in both approaches and not statistically significantly different from one another. We therefore use the estimate from the bucket approach.

(iv) The values for τ_k we use are conservative in the following sense: longer values for τ_k would cause energy use per gsf to rise to higher levels as buildings age, which would mean that the energy savings and NPV of those savings from building and system renewal would be even larger than estimated in the model. Thus the results in the main text can be viewed as lower bounds on the expected savings from renewal of aging buildings and systems, over and above the impact of new technologies that may increase the potential for and reduce the costs of improving energy efficiency in the future.

Having estimated τ_k, it remains to estimate E_k^* in equation (S1). To do so, we set up a panel regression using the full seven-year sample of energy consumption by building and energy carrier:

$$E_k^t = E_k^0 + \dot{E}_k t + aH_t + bC_t + Building\ Fixed\ Effects \tag{S5}$$

where again H and C are heating and cooling degree days for the year. Equation (S5) is similar to (S4) – the difference is that in (S5), we use a panel regression to determine a single estimate for $\dot{E}_k$ by carrier rather than a series of individual estimates for each building, and include fixed effects for buildings.

In eq. (S5) we assume that $\dot{E}_k$ is a constant, that is, we approximate the increase in energy requirements within the estimation period as constant rather than the asymptotic approach to a maximum given by the full formulation in which energy efficiency exponentially approaches a maximum. This allows us to express the slope estimate, $\dot{E}_k$ as:

$$\dot{E}_k = \frac{(E_k^* - E_k^0)}{\tau_k} \tag{S6}$$

The approximation is reasonable given the large estimates of τ_k. Specifically, with τ_k = 117, 37, and 552 years for steam, chilled water and electricity, respectively, using the linear approximation over the 7 year estimation period instead of the exponential function yields errors of approximately 0.2%, 2.0%, and 0.01%, respectively.

The regression intercept E_k^0 and estimates for τ_k determined above allow us to solve for E_k^*. Tables S6-S8 show the full results, including building fixed effects of the panel regression. The estimates for "time" and "intercept" are used to solve for E_k^* as outlined above.

Dependent Variable: Chilled Water per GSF

Variable	Estimate	Standard Error	t Value	Pr > |t|
CS1	2.665655	2.4924	1.07	0.2856
CS2	2.249141	2.4924	0.90	0.3675
CS3	9.130657	2.4924	3.66	0.0003
CS4	2.377287	2.4924	0.95	0.3409
CS5	11.4353	2.4924	4.59	<.0001
CS6	0.186468	2.4924	0.07	0.9404
CS7	10.54404	2.4924	4.23	<.0001
CS8	10.29799	2.4924	4.13	<.0001
CS9	5.861007	2.4924	2.35	0.0193
CS10	1.773518	2.4924	0.71	0.4772
CS11	2.078649	2.4924	0.83	0.4049
CS12	4.163414	2.4924	1.67	0.0958
CS13	10.7412	2.4924	4.31	<.0001
CS14	1.162258	3.0328	0.38	0.7018
CS15	4.50627	2.4924	1.81	0.0715
CS16	3.49626	2.4924	1.40	0.1616

CS17	1.207745	2.4924	0.48	0.6283
CS18	6.788243	2.4924	2.72	0.0068
CS19	4.576603	2.4924	1.84	0.0672
CS20	6.895887	2.4924	2.77	0.0060
CS21	30.5775	2.4924	12.27	<.0001
CS22	2.146807	2.4924	0.86	0.3897
CS23	3.223238	2.4924	1.29	0.1968
CS24	-1.23113	3.0328	-0.41	0.6850
CS25	5.138785	2.4924	2.06	0.0400
CS26	2.053204	2.4924	0.82	0.4106
CS27	3.661614	2.4924	1.47	0.1427
CS28	12.30846	2.4924	4.94	<.0001
CS29	6.752057	2.4924	2.71	0.0071
CS30	5.865578	2.4924	2.35	0.0192
CS31	10.7976	2.4924	4.33	<.0001
CS32	1.045911	2.4924	0.42	0.6750
CS33	1.716594	2.4924	0.69	0.4915
CS34	4.260093	2.4924	1.71	0.0883
CS35	3.883198	2.4924	1.56	0.1202
CS36	7.993075	2.4924	3.21	0.0015
CS37	9.688993	2.4924	3.89	0.0001
CS38	2.70966	2.4924	1.09	0.2777
CS39	1.505326	2.4924	0.60	0.5463
CS40	13.81483	2.4924	5.54	<.0001
CS41	10.42481	2.4924	4.18	<.0001
CS42	10.77078	2.4924	4.32	<.0001
CS43	9.133912	2.4924	3.66	0.0003
CS44	1.385264	2.4924	0.56	0.5787
CS45	1.599566	2.4924	0.64	0.5215
CS46	0.62251	2.4924	0.25	0.8029
CS47	1.124756	2.4924	0.45	0.6521
CS48	0.88007	2.4924	0.35	0.7242
CS49	1.946659	2.6628	0.73	0.4653
CS50	4.665996	2.4924	1.87	0.0621
CS51	1.689892	2.6628	0.63	0.5261
CS52	5.285651	2.4924	2.12	0.0347
CS53	6.445055	2.8052	2.30	0.0222
CS54	0.405419	2.5652	0.16	0.8745
CS55	1.280461	2.4924	0.51	0.6078
CS56	5.479903	2.4924	2.20	0.0286
CS57	3.909796	2.4924	1.57	0.1177
CS58	-0.35226	2.4924	-0.14	0.8877
CS59	16.23978	2.8052	5.79	<.0001
Intercept	-0.12908	3.8560	-0.03	0.9733
Time	0.562459	0.1029	5.47	<.0001
Cooling Degree Days	0.001167	0.00163	0.72	0.4735
Heating Degree Days	-0.00016	0.000500	-0.32	0.7494

Table S6: Regression Results for the Effect of Aging on Chilled Water Consumption

Dependent Variable: Electricity per GSF

Variable	Estimate	Standard Error	t Value	Pr > |t|
CS1	11.24225	3.1537	3.56	0.0004
CS2	10.67439	3.1537	3.38	0.0008
CS3	13.16662	3.1537	4.17	<.0001
CS4	15.14619	3.1537	4.80	<.0001
CS5	42.70161	3.1537	13.54	<.0001
CS6	5.812696	3.1537	1.84	0.0658
CS7	35.26784	3.1537	11.18	<.0001
CS8	-1.39141	3.1537	-0.44	0.6592
CS9	34.35697	3.1537	10.89	<.0001
CS10	10.90218	3.1537	3.46	0.0006
CS11	9.554457	3.1537	3.03	0.0026
CS12	6.413472	3.1537	2.03	0.0424
CS13	10.477	3.1537	3.32	0.0009
CS14	7.642921	3.1537	2.42	0.0157
CS15	9.63675	4.0779	2.36	0.0184
CS16	-1.83323	4.0779	-0.45	0.6532
CS17	6.590108	3.1537	2.09	0.0371
CS18	8.785588	3.1537	2.79	0.0055
CS19	9.662821	3.1537	3.06	0.0023
CS20	23.85107	3.1537	7.56	<.0001
CS21	18.62491	3.1537	5.91	<.0001
CS22	6.966573	3.1537	2.21	0.0275
CS23	108.3127	3.1537	34.34	<.0001
CS24	10.65373	3.1537	3.38	0.0008
CS25	23.63206	3.1537	7.49	<.0001
CS26	20.49013	3.1537	6.50	<.0001
CS27	5.057982	4.0779	1.24	0.2153
CS28	2.290388	3.1537	0.73	0.4680
CS29	11.70701	3.1537	3.71	0.0002
CS30	5.76003	3.1537	1.83	0.0683
CS31	12.61461	3.1537	4.00	<.0001
CS32	13.6256	3.1537	4.32	<.0001
CS33	37.36017	3.1537	11.85	<.0001
CS34	8.767085	3.1537	2.78	0.0056
CS35	35.57098	3.1537	11.28	<.0001
CS36	1.196263	3.1537	0.38	0.7046
CS37	19.27447	3.1537	6.11	<.0001
CS38	43.48276	3.1537	13.79	<.0001
CS39	6.03265	3.1537	1.91	0.0562
CS40	-3.30002	3.7021	-0.89	0.3731
CS41	12.04614	3.1537	3.82	0.0001
CS42	12.48102	3.1537	3.96	<.0001
CS43	0.291853	3.1537	0.09	0.9263
CS44	5.877126	4.7390	1.24	0.2154
CS45	19.59919	3.1537	6.21	<.0001
CS46	30.65492	3.1537	9.72	<.0001
CS47	24.24724	3.1537	7.69	<.0001
CS48	14.75041	3.1537	4.68	<.0001
CS49	6.939184	3.1537	2.20	0.0282
CS50	12.98146	3.1537	4.12	<.0001
CS51	20.25149	3.1537	6.42	<.0001
CS52	24.6033	3.1537	7.80	<.0001

CS53	10.40594	3.1537	3.30	0.0010
CS54	-1.24604	3.4566	-0.36	0.7186
CS55	26.32035	3.1537	8.35	<.0001
CS56	3.84948	3.1537	1.22	0.2227
CS57	6.824995	3.1537	2.16	0.0308
CS58	-4.43291	3.7027	-1.20	0.2317
CS59	3.332467	3.1537	1.06	0.2911
CS60	7.616422	3.1537	2.42	0.0160
CS61	3.17629	3.1537	1.01	0.3143
CS62	2.160422	3.1537	0.69	0.4936
CS63	11.24796	3.1537	3.57	0.0004
CS64	5.862623	3.1537	1.86	0.0635
CS65	4.037998	3.2832	1.23	0.2192
CS66	-4.46023	3.1537	-1.41	0.1578
CS67	14.95334	3.1537	4.74	<.0001
CS68	9.879561	3.1537	3.13	0.0018
CS69	0.917973	3.1537	0.29	0.7711
CS70	-1.60112	3.1537	-0.51	0.6119
CS71	10.41291	3.2832	3.17	0.0016
CS72	1.177486	3.1537	0.37	0.7090
CS73	36.06271	3.1537	11.43	<.0001
CS74	12.7978	3.1537	4.06	<.0001
CS75	5.78581	3.1537	1.83	0.0670
CS76	29.03758	3.1537	9.21	<.0001
CS77	53.02555	3.1537	16.81	<.0001
CS78	20.26105	3.1537	6.42	<.0001
CS79	1.220325	3.2831	0.37	0.7102
CS80	1.675404	3.1537	0.53	0.5954
CS81	-0.5567	3.1537	-0.18	0.8599
CS82	8.909329	4.0779	2.18	0.0293
CS83	0.262534	3.1537	0.08	0.9337
CS84	15.23051	3.1537	4.83	<.0001
CS85	0.406071	3.1537	0.13	0.8976
CS86	3.622785	3.1537	1.15	0.2511
CS87	10.67596	3.1537	3.39	0.0008
CS88	-0.01426	3.1537	-0.00	0.9964
CS89	18.05546	3.1537	5.73	<.0001
CS90	3.227505	3.1537	1.02	0.3065
CS91	13.7568	3.1537	4.36	<.0001
CS92	9.686196	3.1537	3.07	0.0022
CS93	10.9319	3.1537	3.47	0.0006
CS94	4.860957	3.7027	1.31	0.1897
CS95	-0.94061	3.1537	-0.30	0.7656
CS96	-2.56889	3.1537	-0.81	0.4156
CS97	0.290659	3.1537	0.09	0.9266
CS98	-2.62249	3.1537	-0.83	0.4060
CS99	6.026005	3.1537	1.91	0.0565
CS100	1.583125	3.1537	0.50	0.6159
CS101	5.763337	3.1537	1.83	0.0681
CS102	2.913542	3.1537	0.92	0.3559
CS103	1.030993	3.1537	0.33	0.7438
CS104	5.889736	3.7027	1.59	0.1122
CS105	1.710406	3.1537	0.54	0.5878
CS106	2.50427	3.1537	0.79	0.4275
CS107	0.612331	3.1537	0.19	0.8461

CS108	-4.08382	4.7390	-0.86	0.3892
CS109	10.84632	3.1537	3.44	0.0006
CS110	70.01479	3.1537	22.20	<.0001
CS111	15.35357	3.1537	4.87	<.0001
Intercept	2.944208	4.1628	0.71	0.4797
Time	0.460556	0.1119	4.12	<.0001
Heating Degree Days	0.000309	0.000546	0.57	0.5719
Cooling Degree Days	-0.00037	0.00177	-0.21	0.8369

Table S7: Regression Results for the Effect of Aging on Electricity Consumption

Dependent Variable: Steam per GSF

Variable	**Estimate**	**Standard Error**	**t Value**	**Pr > |t|**
CS1	0.018101	0.0195	0.93	0.3531
CS2	0.007152	0.0195	0.37	0.7136
CS3	0.003032	0.0195	0.16	0.8763
CS4	0.007852	0.0195	0.40	0.6869
CS5	0.119792	0.0195	6.15	<.0001
CS6	0.006201	0.0195	0.32	0.7503
CS7	0.204488	0.0195	10.50	<.0001
CS8	0.046779	0.0195	2.40	0.0167
CS9	0.135924	0.0195	6.98	<.0001
CS10	0.007358	0.0195	0.38	0.7057
CS11	-0.03095	0.0195	-1.59	0.1127
CS12	0.00924	0.0195	0.47	0.6354
CS13	0.007777	0.0195	0.40	0.6898
CS14	0.016357	0.0195	0.84	0.4013
CS15	-0.04752	0.0252	-1.89	0.0599
CS16	0.010594	0.0195	0.54	0.5867
CS17	0.019197	0.0195	0.99	0.3247
CS18	0.009668	0.0195	0.50	0.6197
CS19	-0.03392	0.0195	-1.74	0.0821
CS20	0.03429	0.0195	1.76	0.0789
CS21	-0.01686	0.0195	-0.87	0.3870
CS22	0.578426	0.0195	29.71	<.0001
CS23	0.01783	0.0195	0.92	0.3603
CS24	0.017021	0.0195	0.87	0.3825
CS25	0.015879	0.0195	0.82	0.4152
CS26	0.026601	0.0252	1.06	0.2915
CS27	0.070576	0.0195	3.62	0.0003
CS28	0.017349	0.0195	0.89	0.3734
CS29	-0.00206	0.0195	-0.11	0.9156
CS30	0.002829	0.0195	0.15	0.8845
CS31	0.214397	0.0195	11.01	<.0001
CS32	0.034909	0.0195	1.79	0.0737
CS33	0.048061	0.0195	2.47	0.0139
CS34	-0.0057	0.0195	-0.29	0.7697
CS35	0.063785	0.0195	3.28	0.0011
CS36	0.133321	0.0195	6.85	<.0001

CS37	0.020688	0.0195	1.06	0.2886
CS38	0.018079	0.0195	0.93	0.3536
CS39	0.007527	0.0195	0.39	0.6992
CS40	-0.00273	0.0195	-0.14	0.8888
CS41	0.013028	0.0293	0.44	0.6565
CS42	0.007197	0.0195	0.37	0.7119
CS43	0.214486	0.0195	11.02	<.0001
CS44	0.096246	0.0195	4.94	<.0001
CS45	0.029084	0.0195	1.49	0.1359
CS46	-0.00378	0.0195	-0.19	0.8464
CS47	-0.02111	0.0195	-1.08	0.2788
CS48	-0.03161	0.0195	-1.62	0.1052
CS49	-0.04792	0.0195	-2.46	0.0142
CS50	-0.02726	0.0195	-1.40	0.1622
CS51	-0.04451	0.0195	-2.29	0.0227
CS52	-0.00497	0.0195	-0.26	0.7987
CS53	0.033944	0.0195	1.74	0.0820
CS54	-0.00507	0.0195	-0.26	0.7946
CS55	0.044338	0.0195	2.28	0.0232
CS56	-0.03934	0.0195	-2.02	0.0439
CS57	0.007541	0.0195	0.39	0.6987
CS58	0.007603	0.0195	0.39	0.6964
CS59	0.007362	0.0195	0.38	0.7055
CS60	-0.02205	0.0203	-1.09	0.2773
CS61	-0.00044	0.0195	-0.02	0.9818
CS62	-0.00033	0.0195	-0.02	0.9864
CS63	0.003586	0.0195	0.18	0.8540
CS64	-0.0024	0.0195	-0.12	0.9019
CS65	-0.00031	0.0195	-0.02	0.9873
CS66	-0.01053	0.0195	-0.54	0.5890
CS67	-0.00094	0.0195	-0.05	0.9616
CS68	0.006664	0.0195	0.34	0.7323
CS69	0.00554	0.0195	0.28	0.7761
CS70	0.003203	0.0195	0.16	0.8694
CS71	0.138563	0.0229	6.06	<.0001
CS72	-0.00034	0.0195	-0.02	0.9862
CS73	-0.00035	0.0195	-0.02	0.9855
CS74	-0.00035	0.0195	-0.02	0.9856
CS75	-0.00088	0.0195	-0.05	0.9640
CS76	-0.00023	0.0195	-0.01	0.9905
CS77	-0.00035	0.0195	-0.02	0.9856
CS78	0.004769	0.0229	0.21	0.8349
CS79	-0.00034	0.0195	-0.02	0.9861
Intercept	-0.03386	0.0289	-1.17	0.2425
Time	0.005818	0.000802	7.25	<.0001
Heating Degree Days	0.000016	3.952E-6	4.02	<.0001
Cooling Degree Days	2.961E-6	0.000013	0.23	0.8178

Table S8: Regression Results for the Effect of Aging on Steam Consumption

Table S9 summarizes the final parameters used in the model.

	τ_k (years)	E*	E^0	E^{2000}	E^{2005}	Potential Savings (E^{2005}– E^0)
Chilled Water (ton-hrs/yr/gsf)	37	20.8	3.35	5.94	7.48	4.13 (55%)
Steam (klb/yr/gsf)	117	0.675	0.0695	.0950	0.119	0.050 (41%)
Electricity (kWh/yr/gsf)	552	254	14.3	17.8	18.58	4.31 (23%)

Table S9: Final parameters used for the energy model

Conceptually, these parameters have the following meaning. The energy requirements of buildings are assumed to grow at a decreasing rate as buildings and their systems age, following an exponential adjustment to an asymptotic level where τ_k is the time constant for each energy carrier and E^* is the asymptotic energy requirement per gsf. E^* is larger than the actual value at the start of the simulation, E_{2005}. To calculate E_0, we assume that 10 years of savings are available using the assumed formula. Conceptually, E_0 is the minimum energy usage of buildings if all systems were renewed and defects removed. The table compares E_0 to E_{2000} (the earliest year of data) for point of reference. E_0 gives potential savings of 55% for chilled water, 23% for electricity, and 41% for steam. These savings represent the reduction that would be achieved if every defect were repaired and every building system were renewed. As discussed in the paper, these estimates are likely conservative because (i) actual savings from some new construction and renewal projects on campus and in comparable buildings elsewhere have been larger and (ii) the regression estimates omit future technical progress.

The potential energy savings must also be allocated among renewal items and defect categories. As no quantitative data are available, we used the expert judgment of experienced personnel in the facilities and R&M organization to estimate these allocations Table S10 shows how potential savings are allocated among categories. For example, 25% of chilled water savings can be realized by fixing or renewing exterior structures (e.g. repairing or replacing windows), and the remaining 75% can be realized through improvements to HVAC systems. We assume that 33% of electricity savings arise from plug loads unaffected by renewal or maintenance, e.g., laptops, monitors, printers, copiers, etc. Again, this is a conservative assumption since procurement policies can be altered to recommend or requires purchase of more efficient electronics, and behavior change can affect their duty cycle (e.g., turning equipment off when users leave work). These actions, however, are not included in the model. Potential savings within each category are then allocated to individual items in proportion to the renewal cost.

	CW	Electricity	Steam
Exterior Structures	25%	0	25%
Interior Structures	0	0	0
Plumbing	0	0	0
HVAC	75%	33%	75%
Electrical	0	33%	0
Not Renewal Related (e.g. Plug loads)	0	33%	0

Table S10: Allocating Potential Energy Savings among building system categories

S5. Description of Model Formulations

The model is implemented in the Vensim simulation software. The model is divided into 18 sectors, or "views." This section describes each sector of the simulation model. The descriptions provide an overview of the structure and function of each sector. The descriptions in the documentation field of each equation provide more detailed information on each variable and parameter.

I. Building Renewal Sector

The Building Renewal sector captures the stock and condition of buildings and the systems within them, and models the process of aging and renewal.

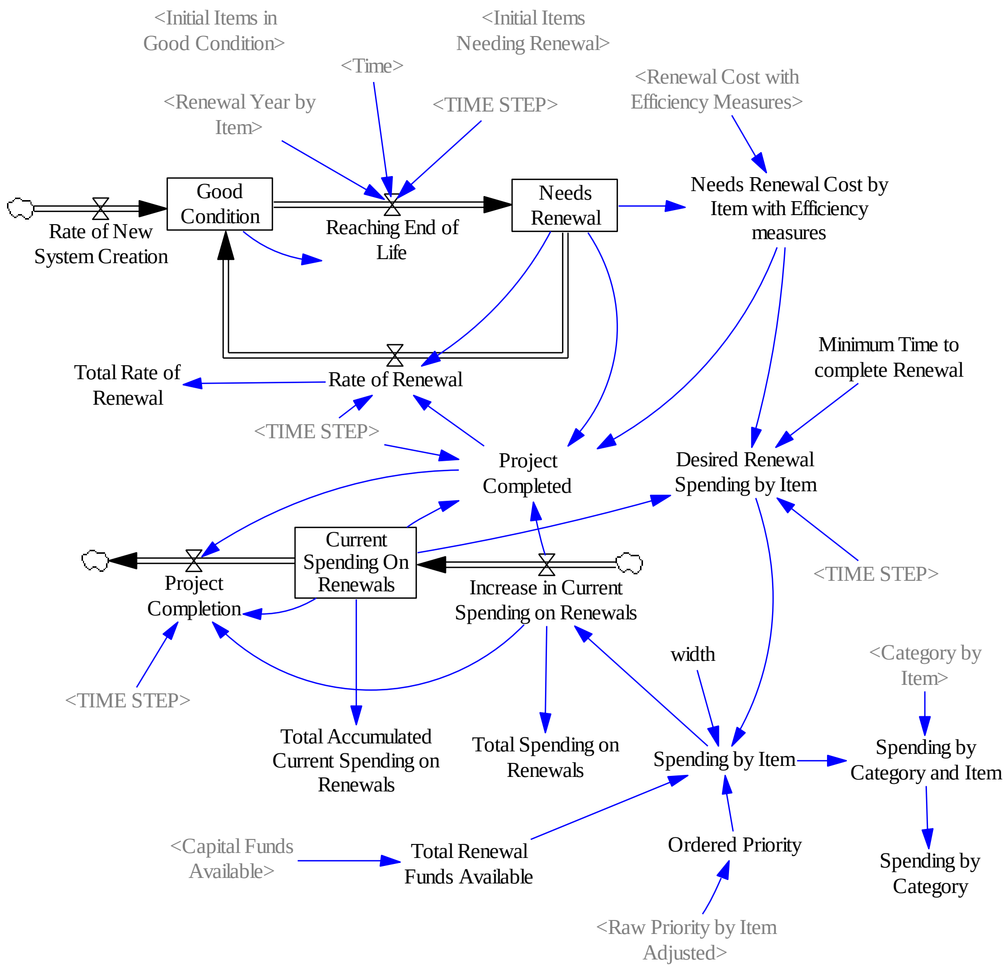

Figure S4: Building Renewal Sector

Following the detailed engineering assessment of buildings and systems described above, each element in the database, roughly 6600 items, is classified as "good condition" (GC) or "needs renewal" (NR). Each individual item is modeled individually in discrete time, moving from the GC stock to the NR stock at the beginning of the year in which it reaches the end of its recommended life. In addition, each item has an associated renewal cost, also provided by the engineering database. The renewal cost is higher if efficiency measures are adopted ("Needs renewal cost by item with efficiency measures"). "Desired renewal spending by item" is the rate of spending required to renew an item. This amount reflects the minimum time needed to complete the renewal, along with any spending that has occurred already. We assume that renewals can be completed in one year if funds are sufficient. Thus, spending is divided evenly over the course of one year, assuming funds are available. The stock "Current spending on renewals" accumulates all spending that has occurred for a given item: when "current spending" matches the "renewal cost" for an item, "project completed" is set to one, the item is moved from the NR stock to the GC stock, and the "current spending" stock for that item is reset to zero.

Total spending is divided amongst individual items according to a prioritization scheme. The variable "ordered priority" assigns a unique rank to each item. The variable "spending by item" uses an algorithm to allocate total renewal dollars among the items. Funds are first allocated to the highest priority item until that item's needs are met, then to the second priority item, and then to the third, and so on until all funds are exhausted. "Spending by category" sums up spending across the categories of building systems.

For the simulations in this paper, prioritization of items needing renewal is random. Random prioritization is a simplifying assumption; in reality, items are completed in groups as entire buildings are renewed. Nevertheless, random prioritization approximates likely future renewal plans for two reasons. First, random assignment with equal probability of selection produces a distribution of renewals across renewal categories (electrical, HVAC, etc.) proportional to the density of items in each category among the buildings, approximating a building-by-building renewal sequence. Second, the order in which particular buildings and systems are selected for renewal are often unknowable and beyond the scope of this analysis. Academic and other programmatic needs, departmental clout, funding availability, donor requirements and unexpected equipment failures can all affect the timing of renewal. Our approach is conservative in that it does not place the items that would yield the greatest energy or maintenance savings at the head of the list, nor does it account for potential synergies and complementarities that would lower the cost of renewal if certain buildings or systems were renewed as a group. For example, it is usually less expensive to renew an entire building and all its systems at once than to do so piecemeal because opening walls for, e.g., plumbing repair also provides access to electrical, HVAC, and other systems. Further, renewing entire buildings or clusters at once can lead to cost savings through use of an integrated design process (Sterman et al. 2014). Similarly, replacing all windows in multiple buildings of similar vintage would lead to economies of scale compared to replacing them in smaller quantities.

II. Capital Spending

The capital spending sector represents the funds available for campus renewal, including additional spending implemented as policies.

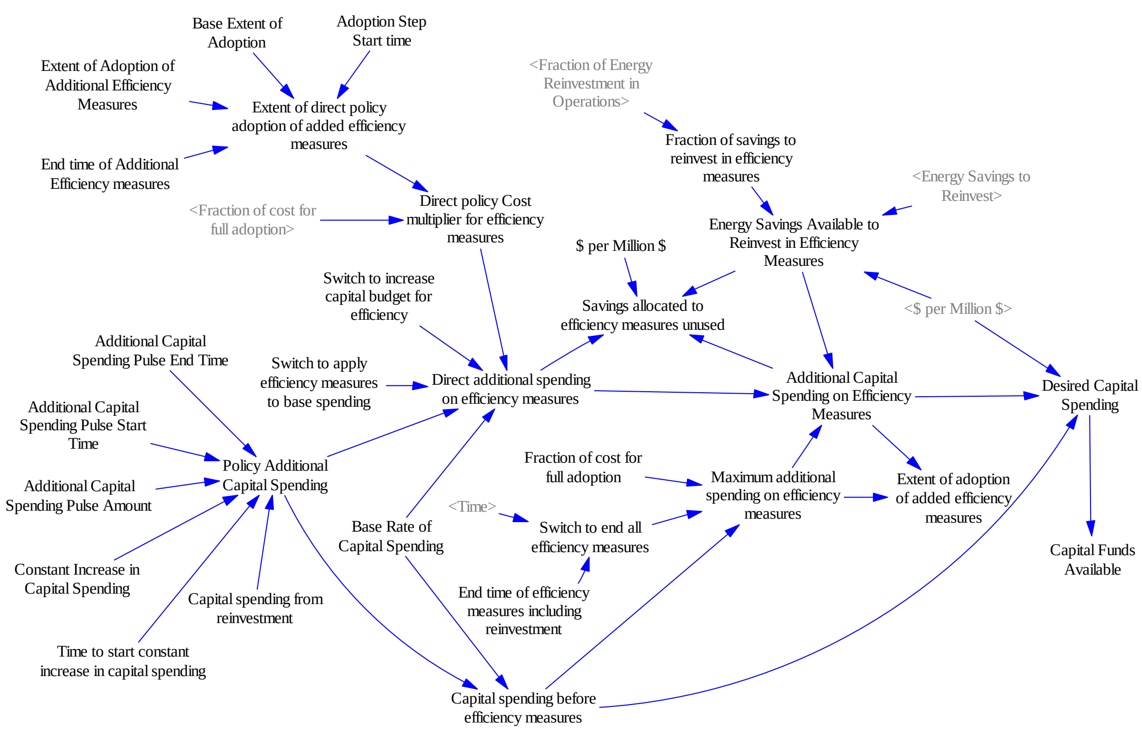

Figure S5: Capital Spending Sector

The amount of capital renewal is a function of capital spending. Capital spending is a policy decision that is featured prominently in the scenarios presented in the main text.

The amount of capital spending is set as an exogenous input. A "base rate of capital spending" is applied in all runs, including "business as usual" (it is set to 19.2M/year, consistent with historical data). "Policy additional capital spending" is the additional spending applied in scenarios. Scenarios can include either a pulse of additional spending with a start and end time, or a constant increase beginning in a specified year.

Capital spending is increased if additional efficiency measures are adopted. The "extent of direct policy adoption of efficiency measures" is a variable between 0 and 1 that reflects a policy choice. The extent of direct policy adoption is used to calculate a multiplier on capital spending ("Direct policy cost multiplier for efficiency measures.")

In addition, efficiency measures can be funded through reinvestment of energy savings from earlier investment ("Energy savings available to reinvest in efficiency measures.") If direct investment in efficiency measures is insufficient, reinvestment makes up the difference. Alternately, savings can be reinvested in maintenance operations, or harvested and spent on programmatic needs.

III. Energy Requirements from Buildings

The next three sectors model the energy demand arising from buildings and their systems. We model the actual energy requirements of buildings in each state (good condition and needs renewal) along with the minimum requirements those buildings would have if they were fully renewed to the greatest efficiency technically possible and the maximum energy use that would be generated if they continued to age without any investment in efficiency.

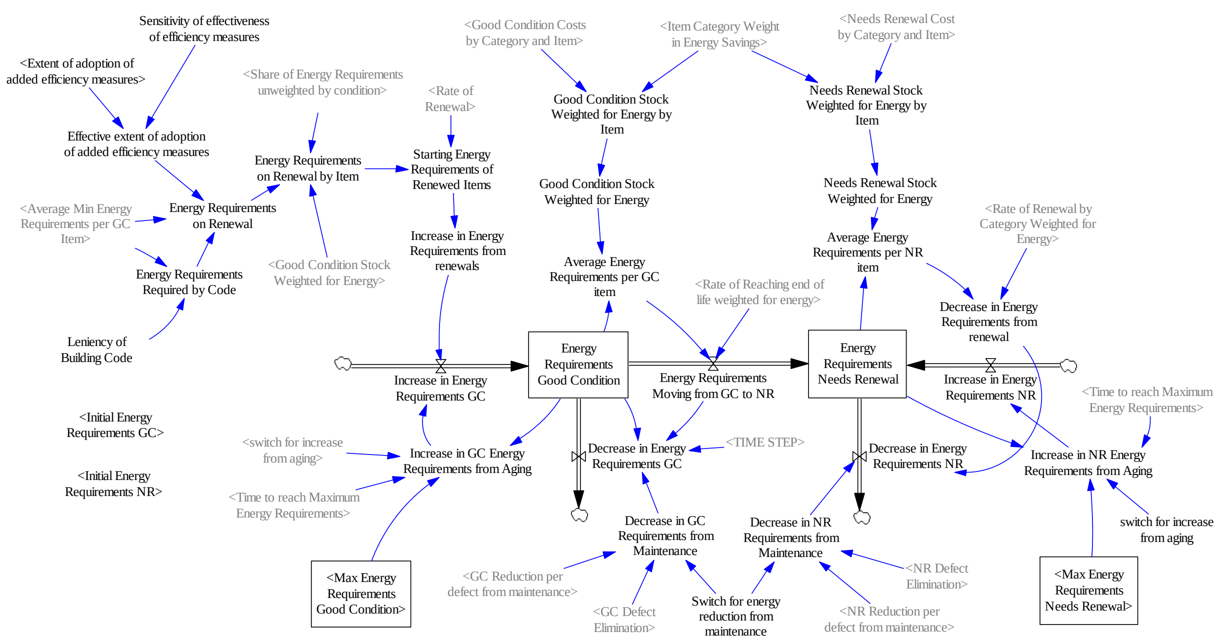

Figure S6: Energy Requirements from Buildings

Buildings have energy requirements (expressed in mBTU/year) that increase with aging and decrease with capital renewal and maintenance.

Energy requirements are modeled using a coflow structure corresponding to the building renewal structure described above. The stocks "Energy Requirements Good Condition" and "Energy Requirements Needs Renewal" track the requirements of building systems in the good condition and needs renewal stocks, respectively. When items reach the end of their recommended life and change state from good condition to needs renewal, the energy requirements of those items change state as well. Similarly, when items are renewed, the energy requirements of those items are debited from the energy requirements NR stock, and the new, updated energy requirements of those items after renewal, which are typically lower, are added to the stock of energy requirements for items in good condition. The energy requirements moving from the good condition state to the needs removal state, and removed from the needs renewal state upon completion of renewal, are given by the average energy requirements for each stock (for details on such "co-flow" formulations, see Sterman 2000).

The two energy requirements stocks are disaggregated into the three energy carriers (chilled water, steam, electricity) and six system categories (exterior structures, interior structures, plumbing, electrical, HVAC, and other). We do not model energy requirements at the individual item level. To calculate the flows of energy demands corresponding to the renewal of individual items, total energy requirements within a category are allocated among individual items based on relative cost and weightings provided by a department of facilities expert.

When items are renewed and re-enter the good condition stock, the energy requirements of the good condition stock increases. Renewal typically improves the energy efficiency of items as, for example, old an inefficient windows, lighting, and HVAC equipment are replaced with modern, more efficient units. The energy requirements of renewed items depend on the technological minimum requirements and the extent of to which efficiency measures are adopted. If adoption of efficiency measures is zero, new items carry the energy requirements required by building codes. As adoption of efficiency measures increases from 0 to 1, the energy requirements of new systems approach the technological minimum. Specifically, the energy requirement on renewal is the weighted average of the minimum requirement and the code requirement, weighted by the adoption of efficiency measures (0-1). Code requirements are defined to be at some level above the minimum, based on the parameter "code leniency."

Figure S6 also shows variables for the increase in energy requirements from aging. Equations and the procedure used to estimate two parameters – the maximum energy requirements and the time to reach the maximum - are described in detail in section S4 above.

Finally, energy requirements are reduced when routine maintenance is done. The calculation of the reduction from maintenance is described below.

IV. Maximum Energy Requirements from Buildings

This sector keeps track of the maximum energy requirements of buildings and systems, which change with the level of efficiency technology embedded in systems as they are renewed.

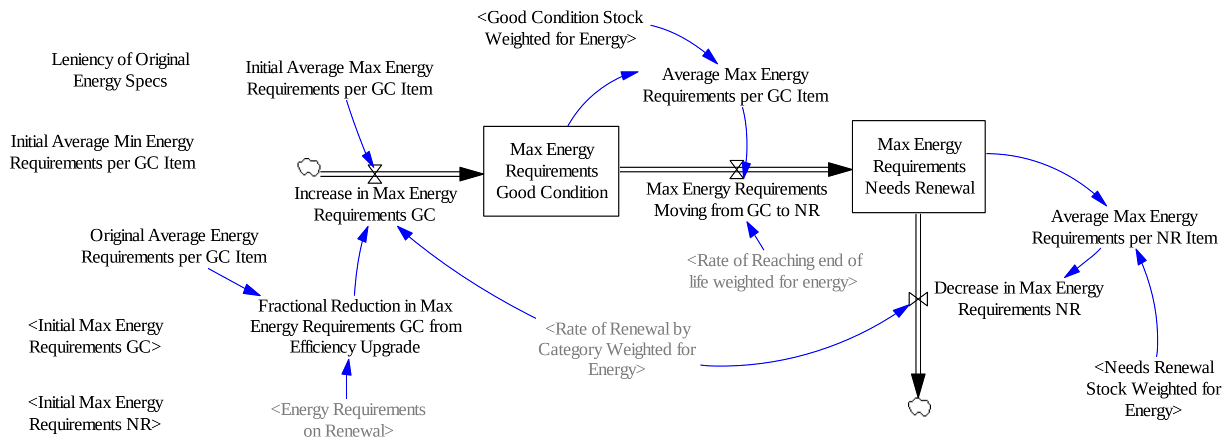

Figure S7: Maximum Energy Requirements from Buildings

A coflow structure is also is used to track the maximum energy requirements. Although the estimation procedure described in section S4 assumes a constant maximum energy requirements per square foot (mBTU/year/gsf), the distribution of total maximum requirements (mBTU/year) between the good condition stock and the needs renewal stock must change as items reach the end of their recommended life and are renewed.

The coflow follows the same principles as above. When items reach the end of their recommended life, the associated maximum energy requirements are moved from the good condition stock to the needs renewal stock. The outflow is calculated by multiplying the rate of reaching the end of life by the average maximum requirements per GC item. Similarly, renewals result in maximum requirements leaving the NR stock.

Our estimation assumes constant maximum energy requirements during the period of estimation (2000-2006). This assumption is reasonable given the low levels of capital renewal and lack of emphasis on efficiency during this time period. In the future, if more efficient building systems are adopted, the maximum requirements of new systems may also change. We assume that more efficient renewals also carry a lower maximum energy requirement, proportional to the ratio between original requirements and new requirements. Thus, the rate at which energy use per item increases as those items age also decreases as more efficient systems are adopted.

V. Minimum Energy Requirements and Reduction in Requirements from Maintenance

This sector keeps track of the minimum energy requirements of buildings and systems technically achievable, which change with the level of efficiency technology embedded in systems as they are renewed.

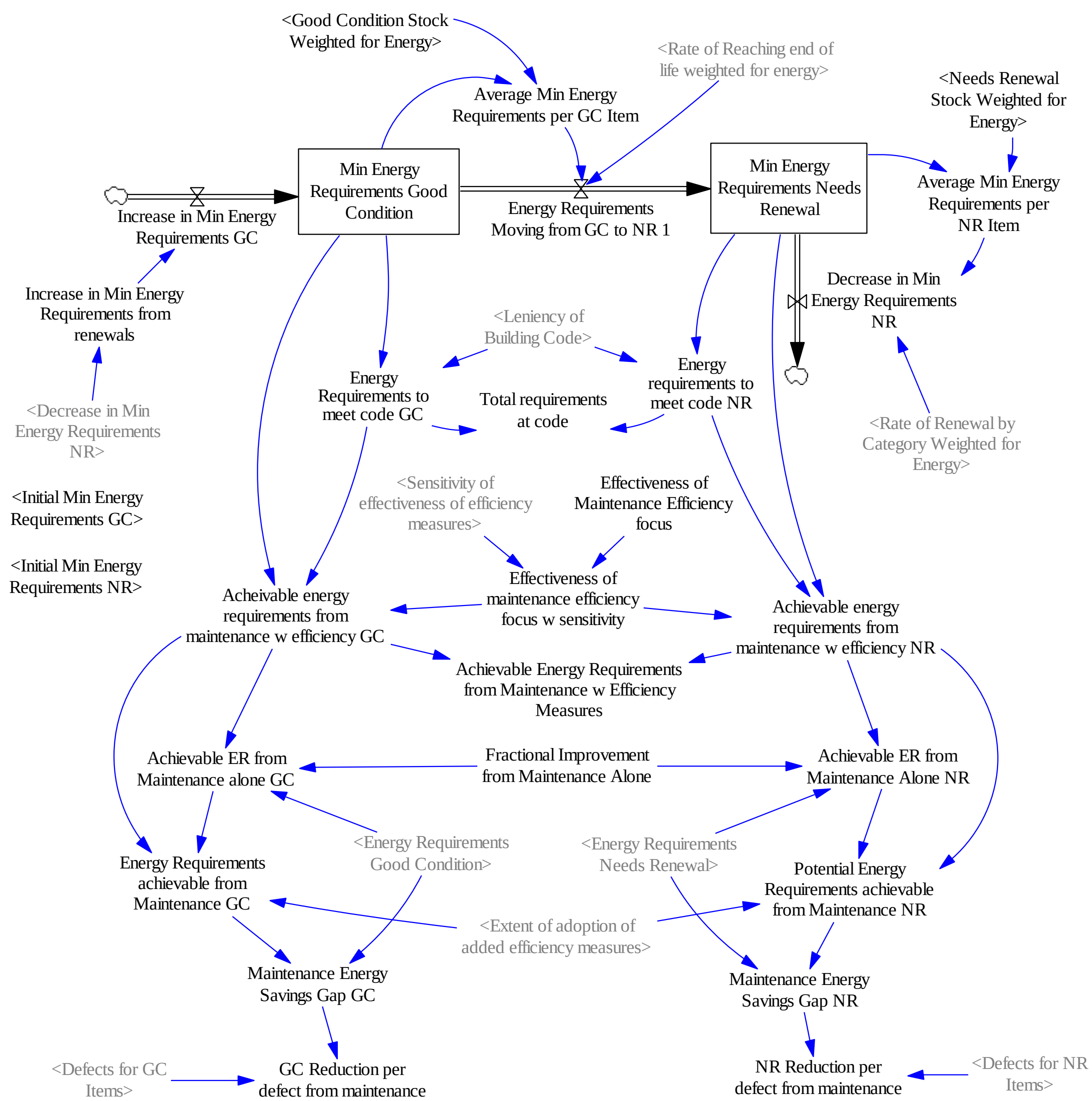

Figure S8: Minimum Energy Requirements and Reductions from Maintenance

A third coflow structure is used to model the minimum energy requirements from buildings. Conceptually, the minimum energy requirement of a building system is the minimum energy consumption achievable using state of the art technology, assuming a constant demand for heating, cooling or electricity. For the purposes of this analysis, we assume a constant minimum requirement per square foot. When items change state from "good condition" to "needs renewal", or are renewed and join the "good condition" stock again, the associated minimum energy requirements of those items also move into the corresponding stocks.

The minimum energy requirements are used to define the code requirements for new installed items. We assume that building codes specifying the least efficient units that may be installed are a multiple of the technically achievable minimum. That is, code requirements are more lenient than the technical state of the art, and are modeled as the minimum requirements multiplied by 1 plus the leniency of code. The more lenient the code, the higher code requirements are relative to the technological minimum.

The minimum energy requirements together with the requirements of building code define the reductions in energy requirements achievable through maintenance. As with capital renewal, maintenance activities can have an efficiency focus. The variable "effectiveness of maintenance efficiency focus" is defined as the extent to which minimum requirements can be achieved when maintenance activities focus on efficiency. If effectiveness is zero, maintenance with efficiency measures brings equipment only back down to the code requirement; if effectiveness is one, maintenance brings efficiency down to the minimum.

The actual reductions achievable from maintenance are a function of the extent of adoption of efficiency measures. If adoption is one, maintenance brings requirements down to the level achievable with an efficiency focus; if adoption is zero, maintenance brings requirements down to the level achievable from maintenance alone. The level achievable from maintenance alone, in turn, is defined as a fractional improvement below current requirements. As an example, when technicians carry out maintenance on fluorescent lighting systems, they can replace old, defective ballasts with similar units or with more efficient models, or replace the entire lighting system with even more efficient LED fixtures and controls, lowering the energy requirements of even further.

Reductions in energy requirements from maintenance are calculated by dividing the total savings available by the number of defects, yielding the savings per defect. When defects are resolved, the rate of defect resolution multiplied by the savings per defect gives the energy savings achieved. These savings are an additional outflow from the energy requirements stocks described above.

VI. Energy Initializations and Totals

This sector initializes the stocks of energy requirements for each category and tracks cumulative energy use.

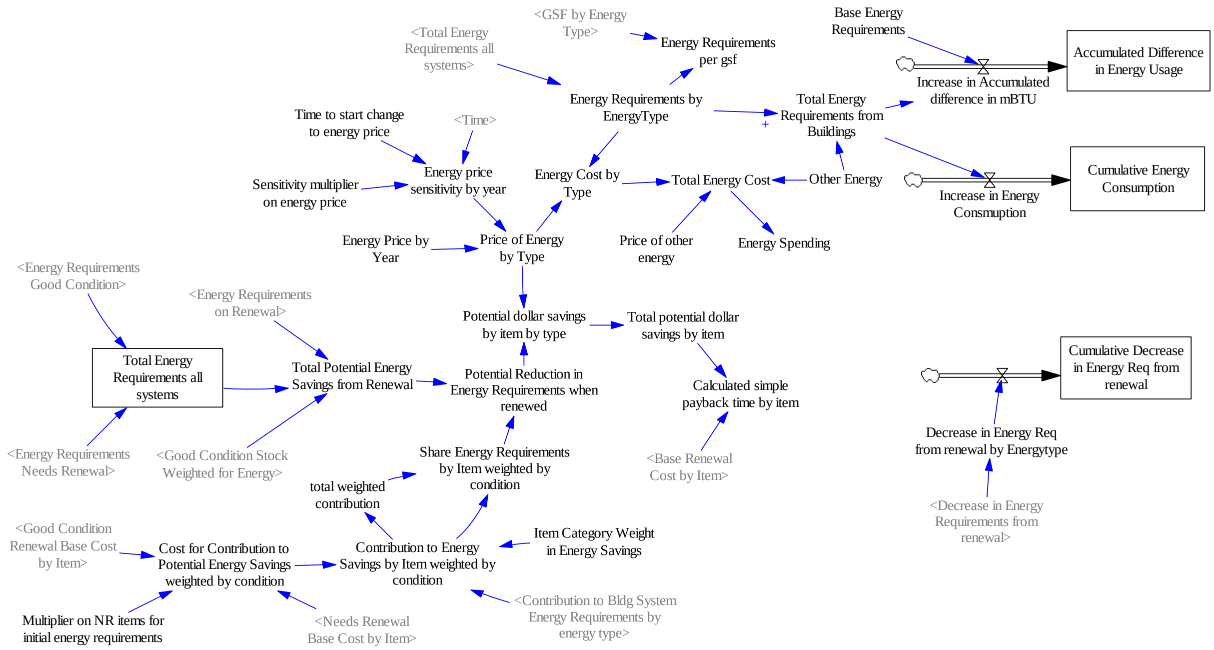

Figure S9: Energy Initializations and Totals

To initialize the stocks of energy requirements for good condition and needs renewal items, we need to allocate total campus requirements (calibrated to match actual campus energy consumption in 2005) among items in the engineering database. The bottom half of Figure S9 shows the factors that figure into this allocation. Requirements are allocated in proportion to the renewal cost of items, weighted in two important ways. First, items past their recommended life are given a higher weight relative to items that are still in good condition (based on the evidence that energy requirements increase as buildings age). Second, items are weighted by the category and type of system, based on expert input (as described in section S4 above). Items that do not contribute to energy consumption, for example, sidewalks, have a weight of zero.

The top half of figure S9 shows energy aggregations and calculations. Annual energy cost ($/year) is given by energy requirements (mBTU/year) multiplied by the energy price ($/mBTU). The energy price varies by carrier, over time, and by scenario. Total energy cost is compared to the base simulation to calculate savings. Savings can, in some scenarios, be reinvested in further energy efficiency and/or additional proactive maintenance. In addition, Figure S9 accumulates energy consumption throughout the simulation, both in total and compared to a base or reference simulation.

VII. Base Defect Creation Rate

This sector models the rate at which defects are created by items in good condition and among those needing renewal. Defects lead to breakdowns, complaints, and other events that generate maintenance work orders.

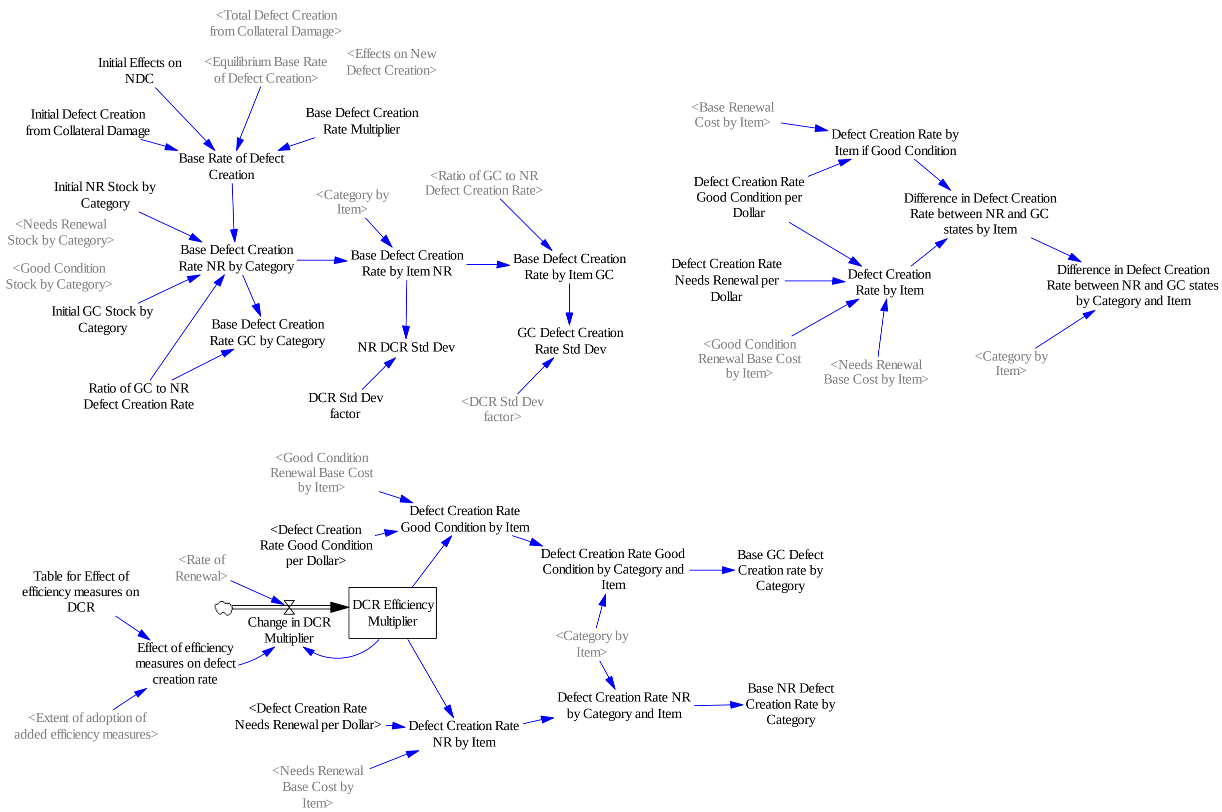

Figure S10: Base Defect Creation Rate

As structures and systems age and wear they generate defects that eventually cause failures, complaints by users or other events that cause maintenance work orders to be opened. The rate of defect creation depends on a number of factors, including whether a system is past its nominal life, the quality of maintenance work, parts quality, and the intensity of use. To calculate the rate of defect creation, we first calculate a "base" rate of defect creation that is a function of system category and whether the system is in good condition or needs renewal. The base rate of defect creation is then adjusted by factors related to maintenance operations (described below).

The base rate of defect creation is the rate of defect creation needed for the stock of defects to be in equilibrium, given the initial distribution of defects between good condition and needs renewal items. (The simulations depart from equilibrium quickly given the rising needs renewal stock). In equilibrium, defect creation must equal defect resolution. We estimate the rate of defect resolution based on the actual rate of maintenance performed per year between 2005 and 2008. (Defect resolution is the rate of work orders closed times the number of defects resolved per work order, a variable described below). Using the rate of defect resolution (defects/year), we then calculate the defect creation rate such that when summed across all items in each building system category the total defect creation rate equals the defect resolution rate. For each item, the defect creation rate is the renewal value ($) multiplied by the density of defect creation per dollar of item value (defects/year/$). The rate is higher when an item is past its recommended life. As a result, as more items age and move into the needs renewal stock, the overall rate of defect creation rises.

Capital investment in additional energy efficiency measures reduces the defect creation rate for some items. State of the art equipment is often more reliable and easier to fix when defects do arise (for example, LED lights last far longer than traditional incandescent, halogen or fluorescent bulbs). A multiplier on the defect creation rate is applied when an item is renewed with adoption of efficiency measures; this multiplier is stored in a stock until the item is renewed again.

VIII. Defect Creation

This sector models the rate of defect creation as it depends on factors including the intensity of system use, collateral damage from breakdowns in other items, part quality and maintenance work quality.

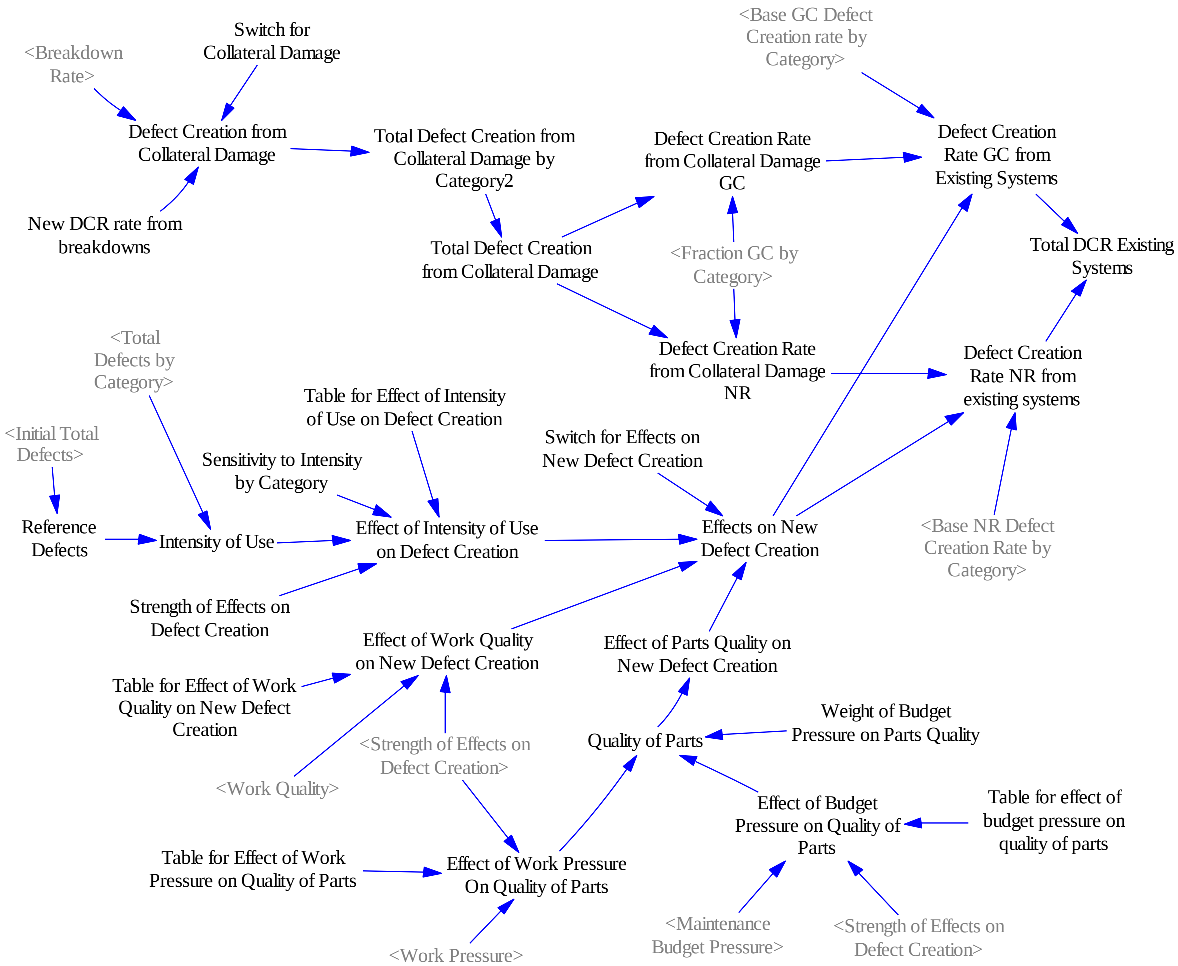

Figure S11 Defect Creation

The actual rate of defect creation is given by the base rate described above, modified by other factors. Defect creation rises when equipment and systems are used more intensively or when maintenance work quality or part quality are low. These effects are multiplied to form the variable "Effects on new defect creation." In addition, maintenance breakdowns can create new defects through collateral damage. For example, the failure of a steam expansion joint can release high-pressure steam into a sub-basement that can damage equipment that was otherwise in good condition. The total defect creation rate is the base rate multiplied by "effects on new defect creation", plus the defect creation rate from collateral damage. The rate is calculated separately for each category, and for good condition items and needs renewal items.

As the defect density rises relative to the reference value, the hazard rate of new defect generation increases. For example, as corrosion compromises a steel column (a defect), the load on nearby columns increases, raising the chance they will fail). The function shown below describes this relationship. The x-axis is intensity of use, and the y-axis is the effect on new defect creation. (The effect saturates and becomes constant beyond the bounds of the graph).

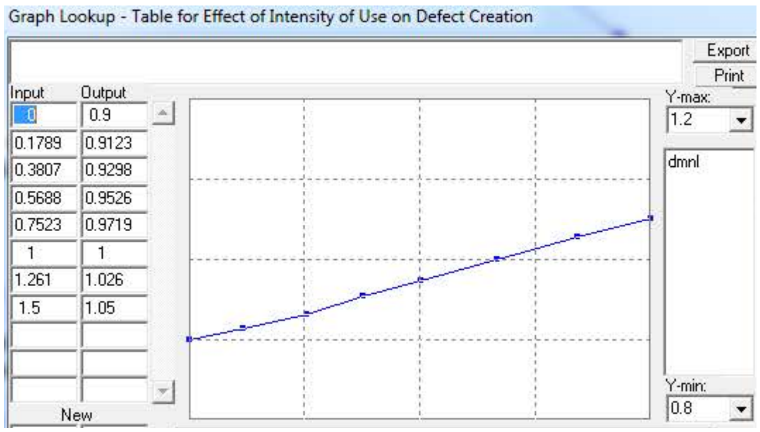

Figure S11: Nonlinear function for the effect of Intensity of use on defect creation

The effects of work quality and parts quality are formulated in a similar manner. Work quality is a function of productivity: As described in the main text, high work pressure causes maintenance technicians to cut corners. Corner cutting boosts productivity (measured as work orders closed per person-hour of effort), but also leads to a decline in the quality of maintenance work. Similarly, parts quality is a function of both work pressure and budget pressure. High budget pressure leads to the use of less expensive, inferior parts; high work pressure reduces the time available to locate parts that are the best match to the need. The variable "Strength of Effects on Defect Creation" moderates all three relationships by adjusting the slope of the functional relationships. A stronger effect implies a higher slope.

Defects from collateral damage are calculated by multiplying the breakdown rate of equipment in each category by a defect creation hazard rate for each breakdown category—defect category combination. The matrix below shows parameter assumptions used, based on expert input from maintenance organization personnel. The total defect creation rate for a defect category is obtained by summing over all breakdown categories.

		Breakdown Category				
		Exterior, Substructure	Interior Structures & Finishes	Plumbing	HVAC	Electrical
Defects Created Category	Exterior, Substructure	0	0	0.05	0.05	0
	Interior Structures & Finishes	0	0	0	0	0
	Plumbing	0	0.05	0.05	0	0
	HVAC	0	0	0.05	0.05	0
	Electrical	0	0	0	0.05	0.05

Table S12: The Collateral Damage Matrix

IX. Equipment Defects and Defect Elimination

This sector keeps track of the stock of defects in items in good condition and among those needing renewal, and models the flows of defect creation and elimination.

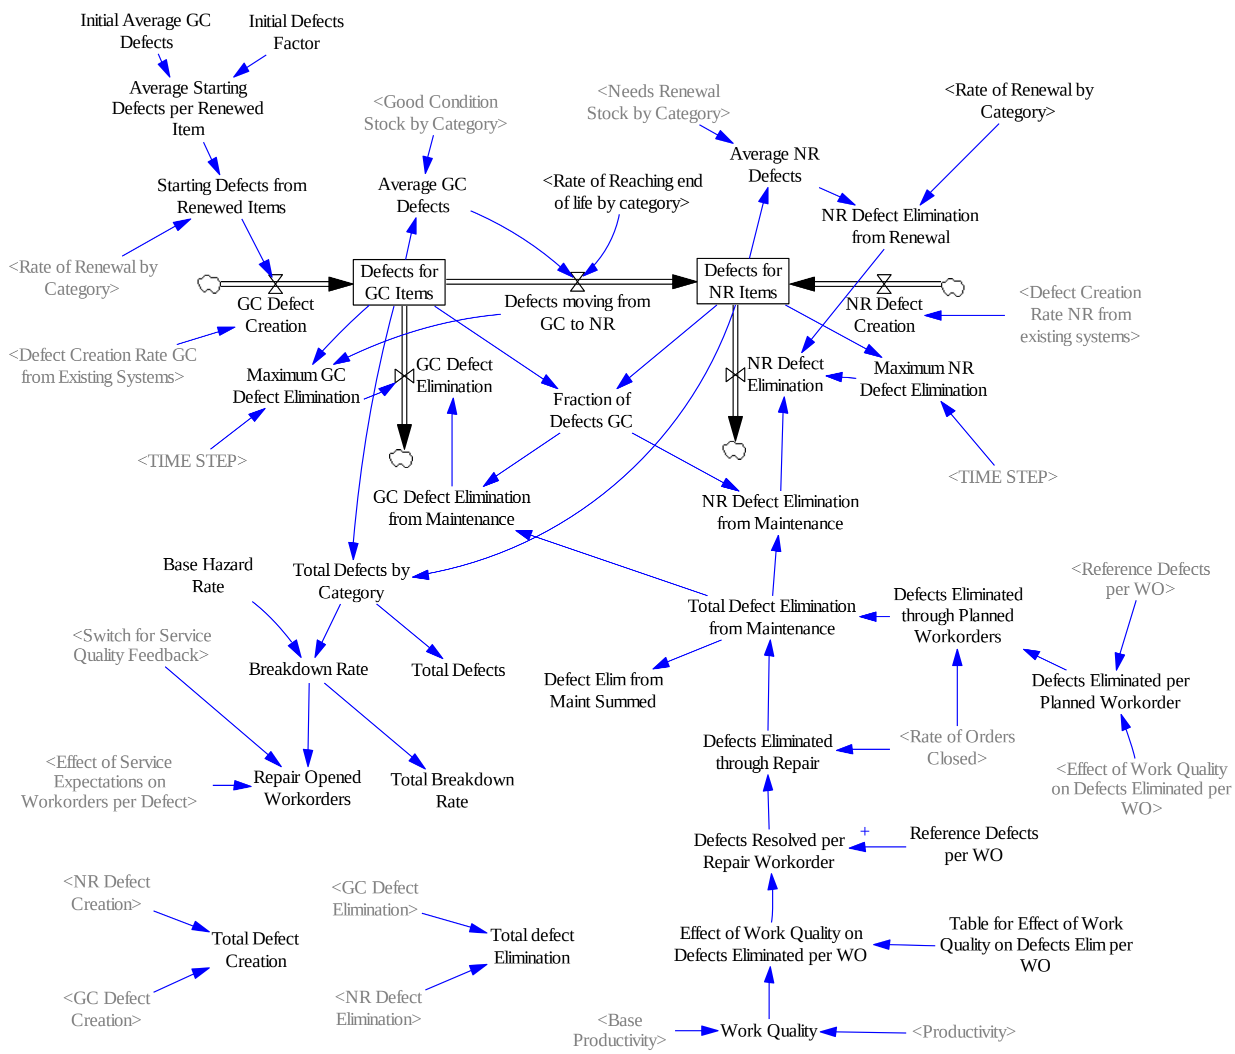

Figure S13: Equipment Defects and Defect Elimination

The model tracks defects in buildings and building systems as they age and are renewed. A defect is defined as a problem that can be reduced through one maintenance work order; thus, large and expensive problems involve multiple defects. Like energy requirements, defects are an attribute of building systems that travel with systems as they reach the end of their lifecycle and move from the good condition to the needs renewal stock. The stock of defects increases as new defects are created and reduced when they are eliminated by maintenance. When items are renewed, the defects associated with them are assumed to be eliminated.

In addition, defects are created through operations, as discussed above, and eliminated by maintenance. Maintenance eliminates defects through two channels: repair (reactive) work and planned (proactive) work. Repair work constitutes responses to breakdowns, and planned work is proactive effort to remove defects before breakdowns occur. In both cases, defect elimination is the number of closed work orders multiplied by the number of defects resolved per work order. In turn, "defects resolved per work order" is a function of maintenance work quality. When work quality is high, more defects are eliminated per work order. As quality slips and workers cut corners, fewer defects are eliminated for each work order.

Figure S13 also shows the determinants of the breakdown rate. The breakdown rate (indexed by category) is equal to the number of defects multiplied by the hazard rate of a breakdown, complaint, or other event that generates a maintenance work order (work orders created per year per defect). The hazard rate is a parameter that varies across categories, as shown in Table S13.

Given the hazard rates of work order generation per defect in Table S13 we can infer the initial stocks of defects by assuming that the stock of work orders is in equilibrium at the start of the simulation. Equilibrium implies that work order creation equals the flow of work orders closed, which we know from the detailed maintenance data provided to us. To estimate the hazard rates of work order generation from defects we note that defects remain latent for different periods across categories. The expected latency period is the expected interval between the time a defect is created and the time it creates a breakdown, user complaint, or other event that results in the generation of a work order. As shown in Table S12, work order hazard rates per defect are lower for exteriors, substructures, interior structures and finishes, and electrical equipment than for HVAC and plumbing systems. The latter involve more mechanical linkages, rotating equipment (motors, pumps, fans), controls, and other components that fail due to aging and wear faster than the components in structures. For example, the failure of a motor, fan belt or fan in an HVAC unit will cause the unit to stop operating, likely causing a user complaint relatively quickly. In contrast, defects such as spalling in the mortar joints of exterior walls may eventually lead to water damage that generates a maintenance work order, but the expected lag between the creation of the defect (cracks in exterior walls) and the work order is far longer. The long residence time before damage and damage detection for many defects provides an opportunity for preventive maintenance: such defects can often be spotted and corrected before they cause breakdowns—if effort is devoted to proactive maintenance.

Category	**Work orders created /year/Defect**
Exterior, Substructure	.06
Interior Structures & Finishes	.07
Plumbing	.134
HVAC	.125
Electrical	.06

Table S13: Hazard Rates for Building System Categories

X. *Work Order Backlog*

This sector tracks backlogs of maintenance work orders and the flows representing the creation and closing of work orders.

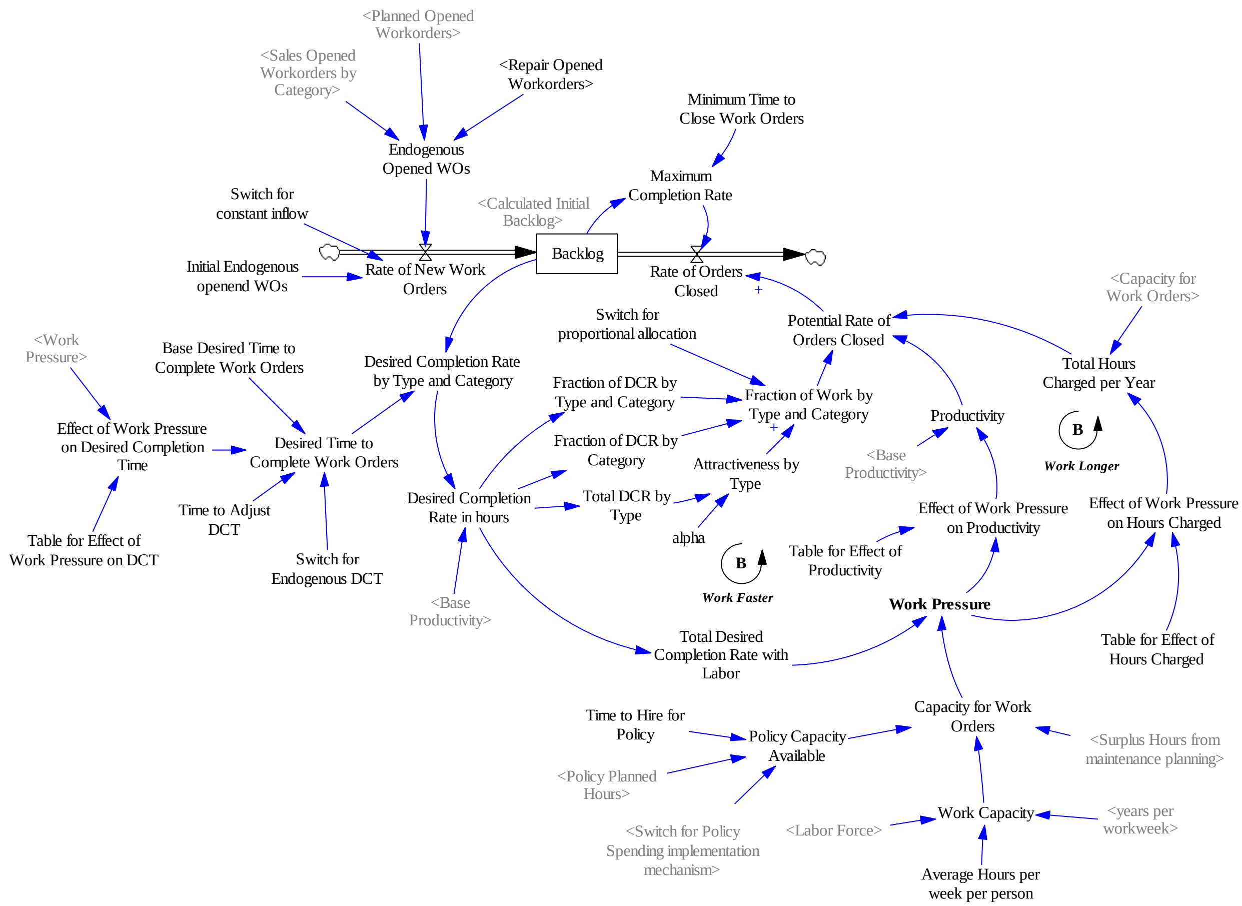

Figure S14: Work Order Backlog and Completion

New maintenance work orders are opened and accumulate in a backlog until they are closed. The model disaggregates the work order backlog by building system category and again by type of work order (proactive or reactive). Work order creation is described above. The flow of work orders closed is determined as follows. We first calculate the desired completion rate (work orders per week) for each type and category of work order by dividing the backlog by the desired completion time. We then calculate desired completion rate in person-hours by dividing by the base productivity of maintenance work in each category. Summing over all types and categories yields the total desired completion rate, in person-hours per week of effort required. The total desired completion rate is then compared to work capacity (person-hours/week of available maintenance labor) to determine "work pressure." The relationship between work pressure and productivity and between work pressure and hours worked is described in the main text. A non-linear function is used in the model (e.g., the "table for effect of productivity) to capture these estimates and saturation limits when work pressure is very high or very low, following the procedure outlined in Sterman (2000, pp. 570-571). The function is linear around the normal operating point, with a slope determined by the regression estimate. The normal operating point is the point at which work pressure =1; that is, where capacity exactly matches the desired completion rate. When work pressure =1, productivity equals "base productivity."

The "fraction of work by type and category" is calculated as follows. We allocate available maintenance resources among categories in proportion to the desired completion rate for each category. Resources are allocated between proactive and reactive work using a logit choice model.

Specifically, the term in eq 2 of the paper s_j, $j \in \{R, P\}$, is the share of maintenance resources allocated to Reactive or Proactive work. These shares are given by a logit choice model,

$$s_j = A_j / \sum_j A_j \tag{S7}$$

$$A_j = \exp(\alpha C_j^*) \tag{S8}$$

where A_j, the affinity or attractiveness of each category and type of work, depends on the desired rate of work completion for each type of work, C_j^*. The logit parameter $\alpha > 0$ was estimated by calibrating the model to data on proactive and reactive work orders opened and closed, hours per work order, workweeks, and backlogs of work orders that determine the desired completion rates. The procedure is equivalent to maximum likelihood estimation where the best estimate of α minimizes the sum of squared errors between the simulated and actual flow of work orders closed, conditioned on the structure for work order backlogs using the actual data for work orders opened and the estimated relationship between work pressure and productivity described in the paper. Finally, the formulation for the completion rate ensures that work orders cannot be completed faster than a maximum rate determined by the backlog and minimum time required to complete work orders. The actual rate of work orders completed is the lesser of this maximum rate and the rate that capacity will support.

XI. Planned Work Orders

This sector represents the rate at which proactive, planned maintenance work is needed and carried out based on the capacity of the maintenance organization.

Figure S15: Planned Work Orders

We endogenously calculate the rate at which planned work orders are opened based on maintenance capacity available. If the demand for reactive work falls below the capacity of the maintenance organization, more planned work orders are opened so that available capacity is fully utilized.

"Capacity for planned work" (measured in person-hours) is equal to mandatary planned hours, plus expected excess capacity in hours, plus any additional increase due to a policy imposed by the model user. Mandatory planned hours are given by scheduled planned maintenance activities that are entered as work orders at regular intervals regardless of workforce capacity; in practice and consistent with the interviews, if reactive demands exceed capacity, these planned maintenance work orders may sit for long durations in the work order backlog. Excess capacity is the total capacity of the workforce, less mandatory planned maintenance, and less the current desired completion rate for reactive work.

"Capacity for planned work" is allocated among the six categories. The rate at which planned work orders are opened for any category cannot exceed the maximum planned work order rate, which is the stock of defects divided by the minimum time to discover defects. Otherwise, hours are allocated in proportion to the desired defect resolution rate by category, which, as discussed above, reflects longer average residence times for structural categories.

XII. Maintenance Staffing

This sector represents the staff level of the maintenance organization, including the desired staff level, hiring, attrition, and layoffs.

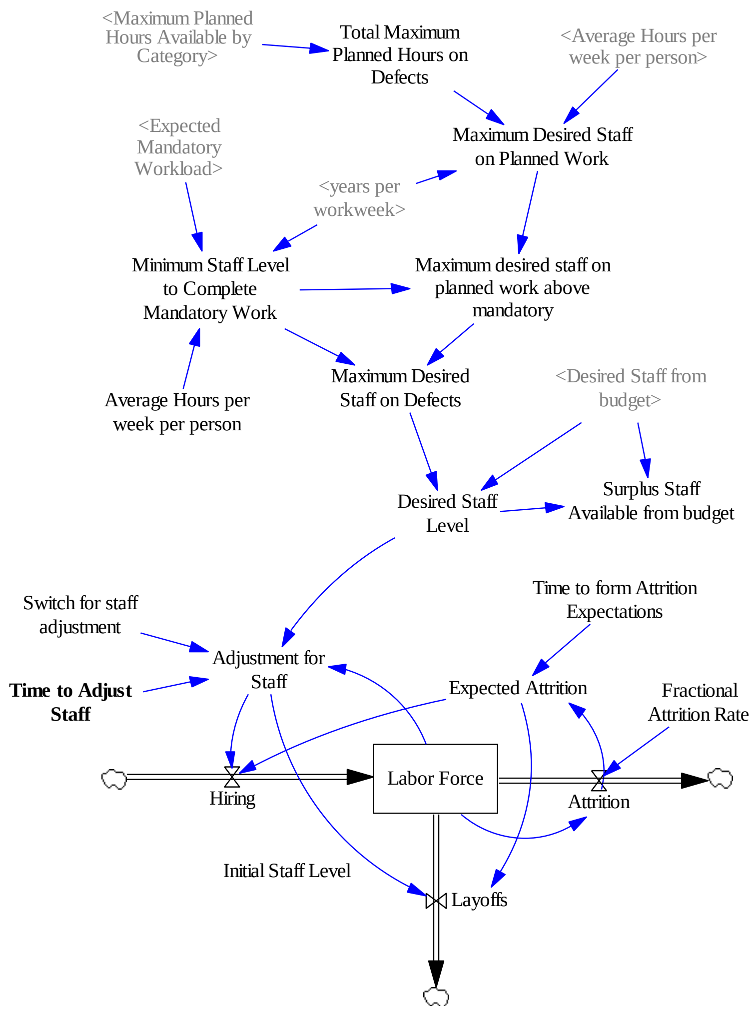

Figure S16: Maintenance Staffing

The maintenance labor force determines the capacity to complete work orders discussed above. The stock is initialized at 100 full time equivalent people. (One FTE person represents 1750 hours per year of effort – 35 hours per week multiplied by 50 weeks per year). The labor force is increased by hiring and decreased by layoffs and attrition.

The variable "Desired Staff Level" determines staff level adjustment. If desired staff is greater than the "Labor Force," the "Adjustment for Staff" is positive and hiring occurs. If the adjustment for staff is negative, layoffs occur. Hiring and layoffs also account for expected attrition.

Desired staff is the lesser of the staff level that can be supported by the current budget ("desired staff from budget") and the staff level that would support all available proactive and reactive work. The formulation ensures that gains in maintenance productivity or reductions in the backlog of deferred maintenance are reinvested. Even if the required reactive work declines the budget and staff level are not cut as long as proactive work remains to be done. The maintenance workforce is reduced only when all available reactive and proactive work is completed. At the start of the simulation, due to the large stock of defects, the maximum desired staff is much greater than the staff level that the budget will support. However, the model is robust under extreme conditions: If a user-implemented policy increases investment in proactive maintenance by a large enough amount, the deferred maintenance backlog can fall low enough that the maximum desired staff begins to fall below the budgeted level, at which point the maintenance staff is gradually reduced through attrition. If the surplus staff is large enough, staff could fall through layoffs; an unlikely situation given the large backlog of deferred maintenance.

XIII. Desired Staff Level from Budget

This sector determines the desired maintenance staff level based on the maintenance budget.

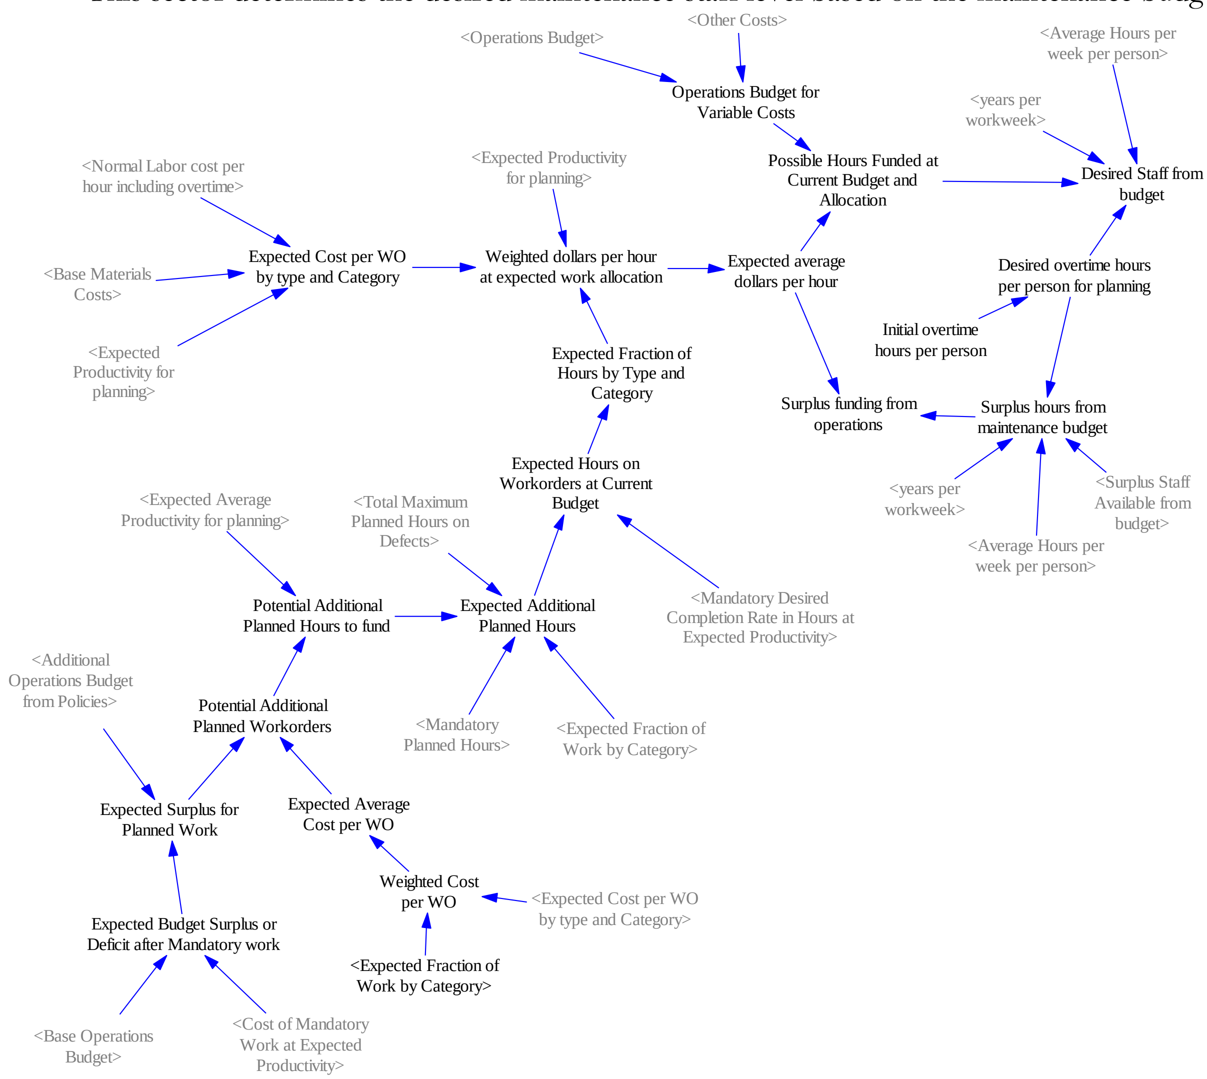

Figure S17: Desired Staff level from Budget

The "desired staff level from budget" is the staff level the current maintenance budget can support, given the expected productivity of the workforce, composition of the work and cost per work order. As average productivity increases, desired staff for a given budget rises. Similarly, if the fraction of work that is proactive rises, due to lower average material costs per hour of work, the desired staff also rises. These factors are captured in the variable "expected average dollars per hour." The maintenance budget divided by expected hours per dollar gives the number of hours of desired capacity. In turn, desired capacity divided by hours per week, multiplied by the workweek gives the desired staff level.

The expected planned fraction of work has an important influence on the number of hires. As a result, for simulations in which maintenance capacity is expanded, we adjust the expected planned fraction in anticipation of the policy change, so that hiring is sufficient to utilize any increased budget.

XIV. Maintenance Budget

This sector represents the maintenance budget and its adjustment to work order demand and investment policies. The maintenance budget funds maintenance operations, and is distinct from capital spending (described above).

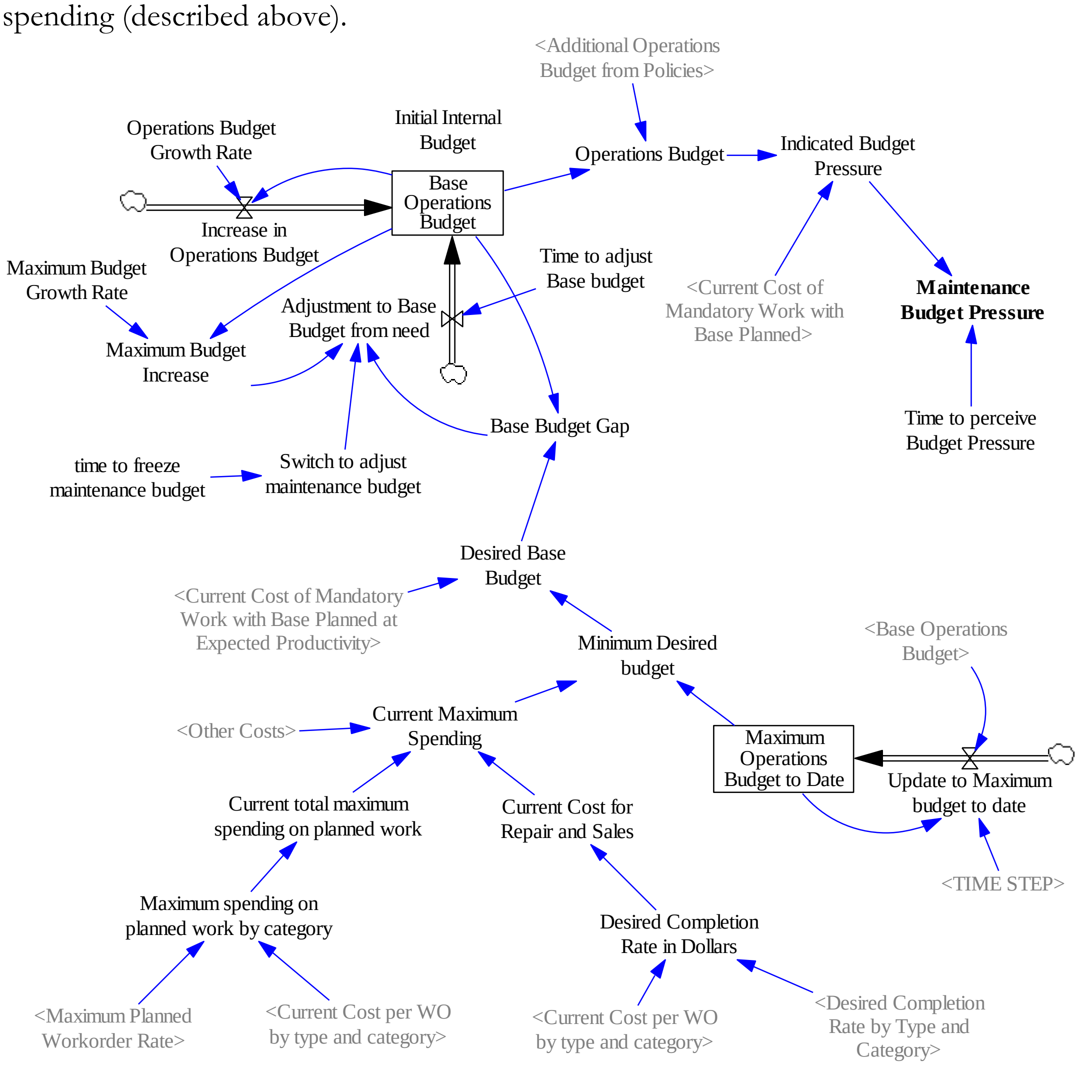

Figure S18: Maintenance Budget

The maintenance budget is a "base maintenance budget" plus any additional budget from policies users may implement. The additional budget from policies reflects specific investment policies as described below and in the main text. The base budget adjusts gradually to a target, or indicated value based on the need to carry out reactive work.

The base operations budget adjusts gradually to match the "desired base budget," based on the parameter "time to adjust base budget." The delay reflects administrative and decision making delays in the budgeting process. If the adjustment time is small, the budget will increase quickly to match increasing reactive work demands. If the adjustment time is longer, staff will not be hired as quickly and reactive work will increasingly crowd out proactive work. A delay time of one year is used, consistent with the budgeting cycle time in the organization. That delay allows the base maintenance budget to grow as the need to carry out reactive maintenance increases, or if the cost of reactive and mandatory planned work rises. However, consistent with the approach used by R&M management, we assume that the desired budget will not fall as long as proactive work still remains to be done. Thus, the desired budget is the maximum of required mandatory spending and a minimum budget floor. If mandatory spending rises above the floor, the desired budget increases. If mandatory spending falls (for example, when proactive investment reduces the volume of reactive work), the desired budget remains at the floor. The budget floor, in turn, is the minimum of past budgets and the spending required to complete all work. Thus, the budget cannot be reduced until proactive work no longer remains.

XV. Maintenance Investment Policies

This sector represents specific maintenance investment policies, including increased spending and reinvestment of energy savings.

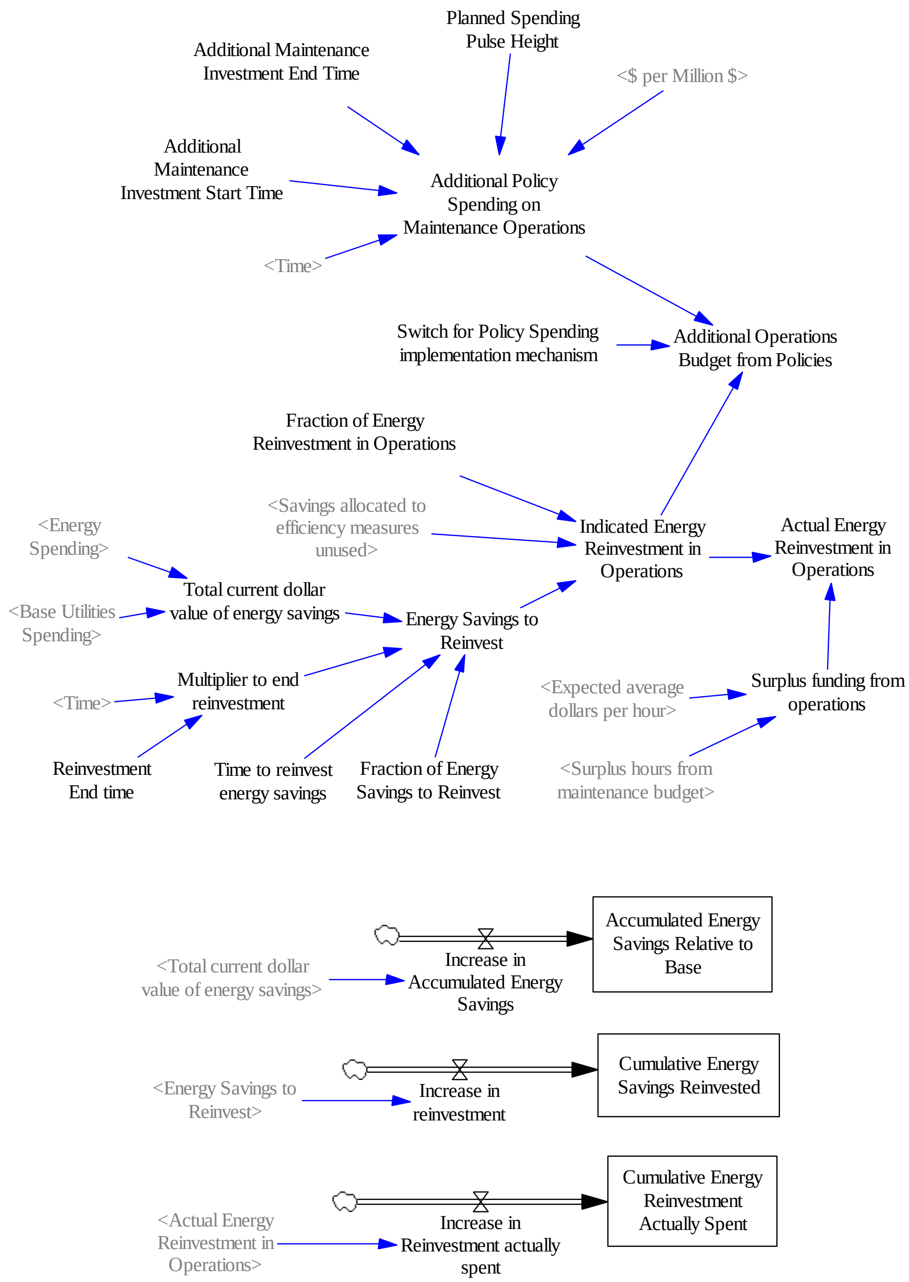

Figure S19: Maintenance Investment Policies

The maintenance budget can be increased by two main policies: (1) direct investment, and (2) reinvestment of savings on energy use. Direct investment is formulated as a pulse where the height, start time and end time are specified by the user. Reinvestment of savings from energy efficiency programs is calculated by subtracting current total energy spending from the level of spending on energy at the corresponding time in the base (or reference) simulation. The formulation estimates the savings resulting from efficiency programs compared to the business as usual situation in which those programs are not implemented. In addition, energy price variations affect energy spending and hence whether there are savings that can be reinvested. The difference is then multiplied by the fraction of energy reinvestment in operations, which captures the extent to which the savings from increased energy efficiency or lower energy prices are in fact reinvested in further improvement or harvested and allocated to other Institute needs.

The model also accumulates the total difference in energy spending over the course of the simulation compared to the reference or comparison simulation, the total amount of savings available for reinvestment, and the total amount actually reinvested, which can be lower than the amount available if the need for additional maintenance spending falls below the available funding.

XVI. Spending and Investment Totals

This sector represents total spending and investment for maintenance operations, capital investment, and energy.

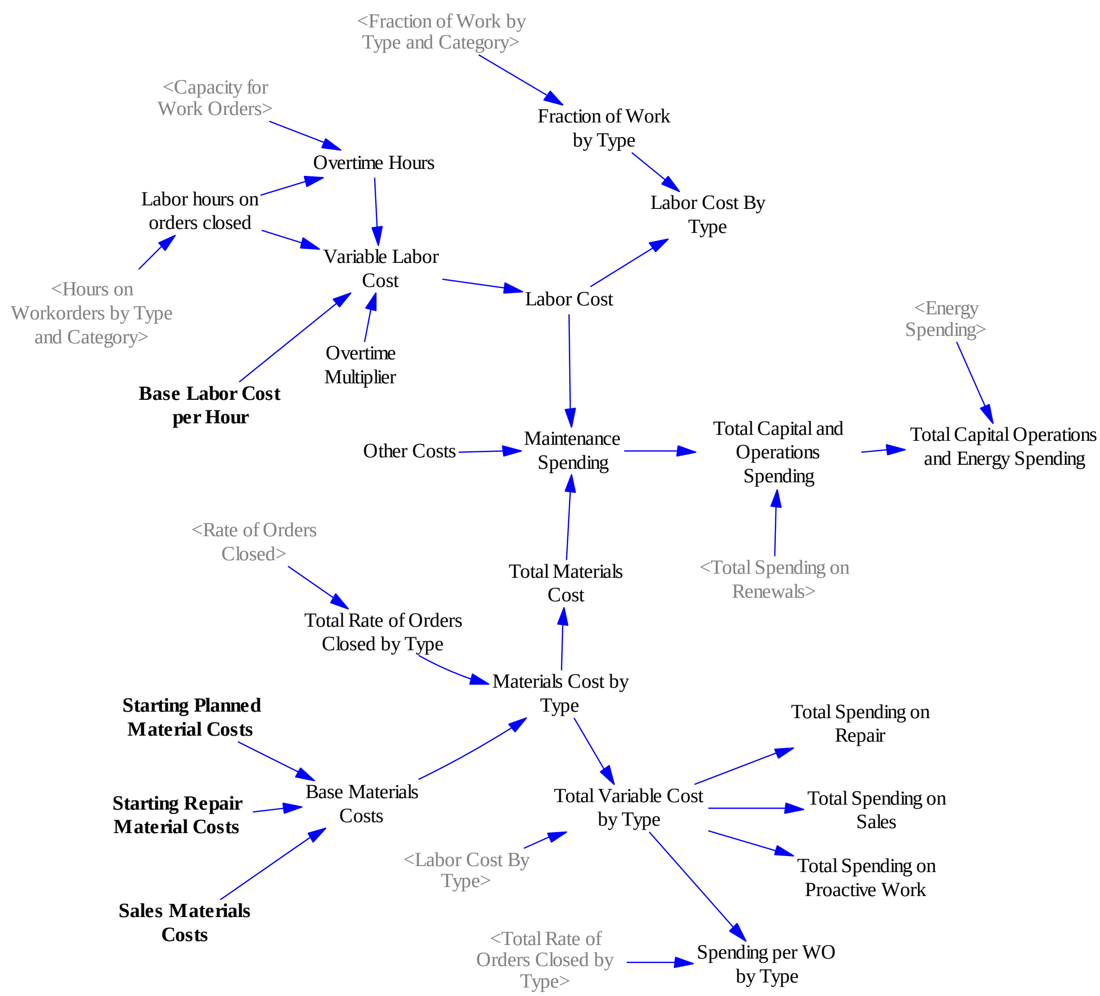

Figure S20: Spending and Investment Accumulations

Maintenance spending has three components: labor, materials, and fixed costs. Labor spending is given by total hours worked and the hourly wage, with time and a half for overtime. Materials spending is the rate of work orders closed multiplied by the cost per work order. The entire model is run in real (2005) dollars, and we assume materials costs and wages remain constant in real terms.

Maintenance spending is added to capital investment to give total capital and operations spending. This result is added to energy spending to give total capital, operations and energy spending. Energy spending is determined by the price of each type of energy and the "energy requirements of buildings" described above. In the base run, we assume that energy prices are constant in real terms. Sensitivity analysis for different energy price trajectories are discussed in the main text.

XVII. NPV of Investment Policies

This sector calculates the net present value (NPV) of investments compared to a reference investment. NPV is calculated for maintenance and operations only (excluding energy), and for maintenance, operations and energy combined.

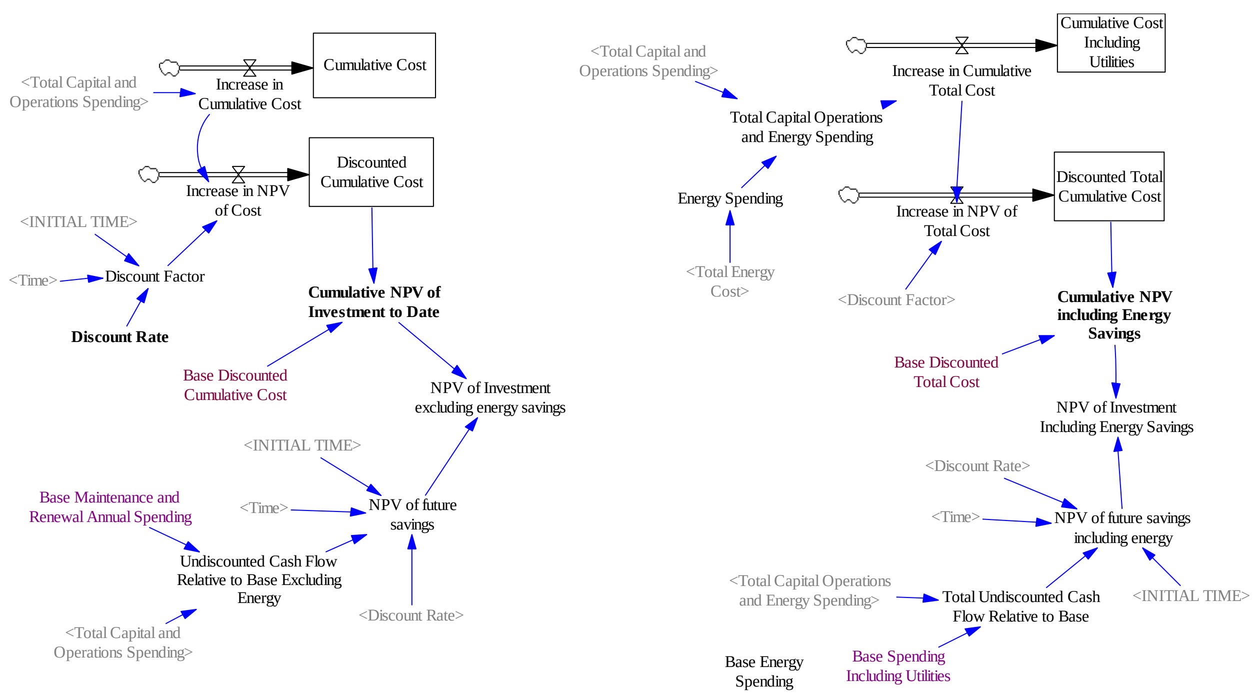

Figure S21: Calculating the NPV of Investment Policies

The NPV of each policy simulation has two components: the difference in discounted cumulative spending through the end of the simulation (in 2030), plus an estimated terminal value of future savings beyond 2030. The discount rate is assumed to be 5% in most simulations, consistent with the actual cost of capital, but is varied in sensitivity testing.

For each simulation, we accumulate discounted spending in a stock of "cumulative discounted spending." At any point in time, discounted spending is undiscounted spending multiplied by a discount factor, exp(-*rt*), where *r* is the discount rate and *t* is elapsed time since the start of the simulation. We store cumulative discounted spending for a base run, and compare the base against the same variable for each policy run. A similar comparison is made for the terminal value. To calculate the terminal value, we assume that the difference between base spending and policy spending in the final step of the simulation remains constant, and discount this perpetuity at the same rate *r*.

XVIII. Calculating Discounted Investment Costs

This sector calculates discounted investment costs. Discounted costs are used to calculate the benefit/cost ratio of investments.

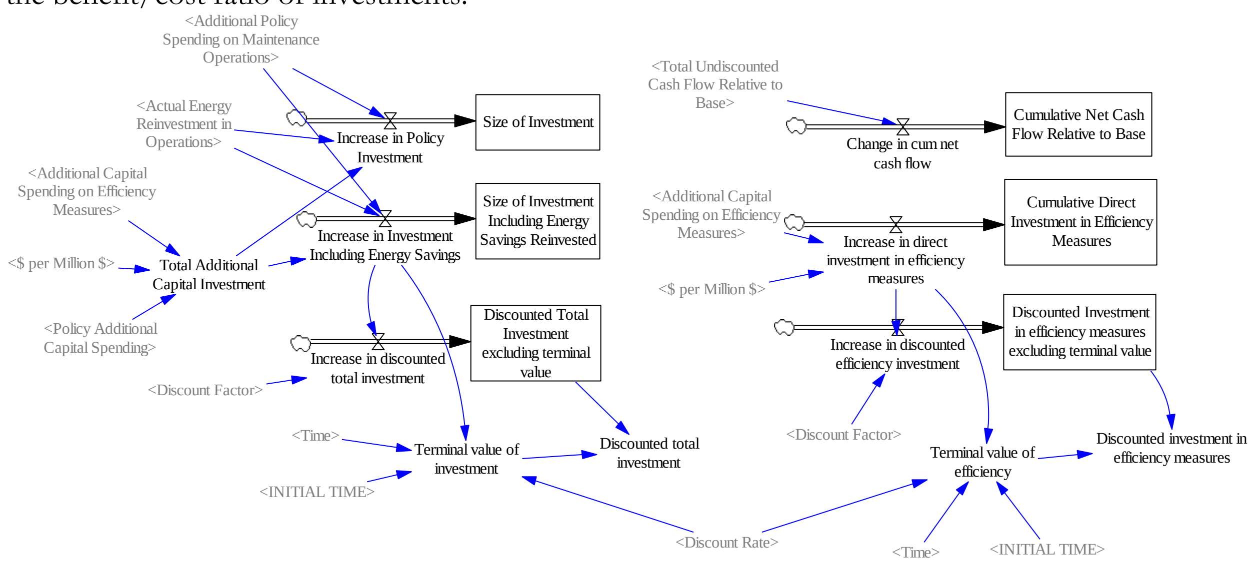

Figure S22: Calculating Discounted Investment Costs

Discounted investment costs are calculated in a manner similar to the NPV calculation described in the previous section. The various rates of investment are discounted by the factor exp(-*rt*), and then accumulated in the stock "Discounted Total investment excluding terminal value". The final investment rate is then used to calculate a terminal value, assuming that the investment will continue at the same rate into perpetuity. The terminal value is then added to the accumulation to yield the discounted total investment.

The process is carried out for both investments in capital renewal and operations, and for investments in additional efficiency measures.

S6. Simulating the Model and Replicating Results

The model is implemented in Vensim, version 5.9. We use Euler integration with a time step of 2^{-6} (0.015625) years, which is short enough to provide a good approximation to the underlying continuous time dynamics of the system (reducing the time step to 2^{-7} years produces no materially significant changes in results).

The model calls an external excel file ("Data for vensim.xls") containing the database of building components and systems, which is used to track the aging of systems from the "good condition" state to the "needs renewal" state. The Excel file also includes item weightings for contributions to energy requirements, energy prices over time by energy carrier (past and future assumptions), and stored results for several variables from a comparison run. The comparison data is used to calculate the difference in energy consumption and cash flow between a current policy and a comparison policy. The difference is needed to calculate energy savings available to reinvest and the NPV of an investment policy relative to a comparison such as Policy 1.

We include the source code (Vensim .mdl file) for the model with this documentation. The full database of building components, conditions, renewal dates and renewal costs is confidential and is not included here. Readers interested in replicating the simulations should contact the authors.

Below we provide the parameter values for all simulations reported in the paper and steps for creating a comparison run.

The steps to create a comparison run are as follows:

1. Under model settings, check the box to save results every TIME STEP. Due to the large number of model variables, for policy runs we save model results every 0.25 years. However, to increase the accuracy of comparisons we store results for each time step for the comparison run.
2. Simulate the model.
3. Export comparison variables, by selecting **Model -> Export dataset** and choosing the name of the simulation just created. Use the file "base run save list.lst" to specify the variables to output. Export to Excel, and select "time running down".
4. Copy the output to the tab "FinancialBase" in the Data for vensim.xls excel file, ensuring that the position and order of variables remains the same.
5. Resimulate the model, to test that the correct data are loaded. The difference in cash flow between the current model and the comparison model should now be zero throughout the simulation.

A comparison run is not necessary for the "Maintenance surge" and "Renewal surge" simulations (Figure 6 in the main text), because these simulations do not involve reinvestment of energy savings or NPV calculations. For all of the policy and runs shown in Figure 7 and Table 1, we use Policy 1 (continuous renewal) as the comparison run.

The parameter changes in Table S14 were used to create the simulations described in the paper. All other parameters retain their base or default values. Running the simulation with no changes (all default values) produces the business as usual (BAU) simulation shown in Figure 6. Variables with changes can be seen on the "dashboard" view in the model.

Note that the "extent of adoption of efficiency measures" refers to direct adoption through increased capital spending. (For example, in Policy 3, additional capital spending is increased from $150M per year to $154.2M per year to fund additional efficiency measures). In Policy 4, actual adoption of efficiency measures reaches 100% as reinvested energy savings are used partly for efficiency measures. In Policy 4, efficiency measures are ended in 2020 after simulations reach a saturation point where the majority of energy-focused renewals have occurred and funds are better harvested.

Model outputs are also displayed on the "dashboard" view. Outputs reported in Table 1 are either the last value (2030) of variables listed on this view, or can be derived as simple calculations using these variables. The discounted benefit cost ratio is derived from the "NPV of Investment Including Energy Savings" and the "Discounted Total Investment." (Benefit cost ratio = 1 + Net NPV/NPV of Investment). Payback time is defined as the number of years after the start of the investment when the variable "Cumulative net cash flow relative to base" is zero.

	Planned Spending Pulse Height (Million $/Year)	Additional Capital Spending Pulse Amount (Million $/Year)	Constant Increase in Capital Spending (Million $/Year)	Extent of Adoption of Additional Efficiency Measures (0 = no adoption, 1 = full adoption)	Fraction of Energy Savings to Reinvest (1 = full investment)	End time of efficiency measures including reinvestment
Business As Usual (BAU)	-	-	-	-	-	-
Maintenance Surge	5	-	-	-	-	-
Renewal Surge	-	150	-	-	-	-
Policy 1: Sustained Renewal	-	-	150	-	-	-
Policy 2: Policy 1 + Maintenance Surge	5	-	150	-	-	-
Policy 3: Policy 2 + Additional Efficiency Investment	5	-	150	0.5	-	-
Policy 4: Policy 3 + Reinvestment of Energy Savings	5	-	150	0.5	1	2020

Table S14: Parameter Changes Needed to Replicate Simulations (Figures 6 & 7, Table 1of paper)

Replicating the sensitivity results displayed in Table 2 requires creating a new comparison run for each condition. (The exception is the energy efficiency potential – for these tests, Policy 1 from above can again be used as a comparison because the renewal only policy involves no efficiency investments). To test sensitivity to the discount rate or to energy prices, Policy 1 is used as the comparison run, with additional changes made to either the discount rate or to the variable "sensitivity multiplier on energy prices."

Policy 4 is then used again with the following changes:

- "Discount rate" set to 0.03 or 0.09
- "Sensitivity multiplier on energy prices" set to 1.5 in the strong condition, and to 0.8 in the weak condition.
- "Sensitivity of effectiveness of efficiency measures" set to 0.75 or 1.1

To provide additional information on the results, Figure S23, below, shows total maintenance spending in the BAU case and policies 1-4. Note how P4, in which energy savings are reinvested, leads to substantially more maintenance spending before 2022, substantially reducing the stock of defects and reactive maintenance costs, leading to still greater savings that further boost both maintenance and energy efficiency – P4 significantly strengthens the reinforcing "Reinvestment" feedbacks compared to the other policies.

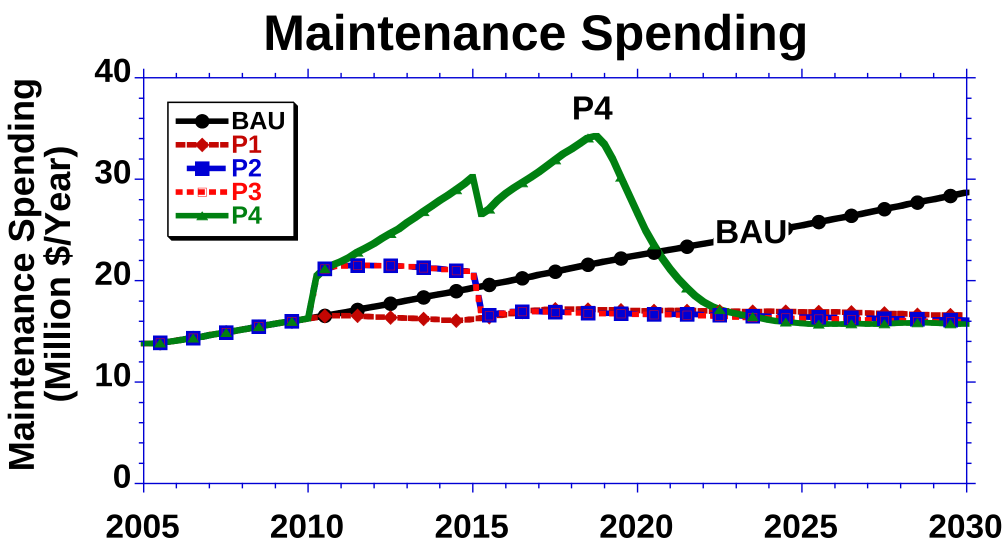

Figure S23. Maintenance spending in the BAU scenario compared to Policies 1-4.